Towards Interactive Recommending in Model-based Collaborative Filtering Systems

Benedikt Loepp
University of Duisburg-Essen
Duisburg, Germany
benedikt.loepp@uni-due.de

Jürgen Ziegler
University of Duisburg-Essen
Duisburg, Germany
juergen.ziegler@uni-due.de

ABSTRACT

Numerous attempts have been made for increasing the interactivity in recommender systems, but the features actually available in today's systems are in most cases limited to rating or re-rating single items. We present a demonstrator that showcases how model-based collaborative filtering recommenders may be enhanced with advanced interaction and preference elicitation mechanisms in a holistic manner. Hereby, we underline that by employing methods we have proposed in the past it becomes possible to easily extend any matrix factorization recommender into a fully interactive, user-controlled system. By presenting and deploying our demonstrator, we aim at gathering further insights, both into how the different mechanisms may be intertwined even more closely, and how interaction behavior and resulting user experience are influenced when users can choose from these mechanisms at their own discretion.

KEYWORDS

Recommender Systems; Matrix Factorization; User Experience

ACM Reference Format:
Benedikt Loepp and Jürgen Ziegler. 2019. Towards Interactive Recommending in Model-based Collaborative Filtering Systems. In *Proceedings of the Thirteenth ACM Conference on Recommender Systems (RecSys '19), September 16–20, 2019, Copenhagen, Denmark.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3298689.3346949

1 INTRODUCTION

Recommender systems (RS) have become very popular in a wide range of application domains, supporting users in finding items that match their interests. The most frequently used method is *collaborative filtering* (CF), which exclusively relies on explicit or implicit feedback provided by the user community for the items. One of the most effective and efficient CF techniques is *matrix factorization* (MF) [7]. When employing a MF algorithm, an abstract model consisting of a number of latent factors is derived from the underlying user-item feedback data. While this leads to very accurate results in terms of objective quality metrics, the possibilities for users to interact with MF recommenders and to influence recommendations are mostly limited to (re-)rating single items. Moreover, as latent factor models are entirely statistical, it is difficult to comprehend the recommendation process. These issues are prevalent in model-based CF in general, although it is long known that aspects related to user experience considerably contribute to user satisfaction [6].

Based on our past research on improving user control in model-based CF [2, 9, 10], we in this paper present a demonstrator that uses the newly implemented *TagMF* framework and for the first time holistically integrates our proposed approaches. This way, we illustrate that it is easily possible to extend the typically fully automated contemporary matrix factorization RS with interactive techniques, thus overcoming several of the widely discussed drawbacks of this kind of method. Our objective is to gain further insights into how (our interactive but also other) recommendation components may be combined with each other more closely, and to offer a means for future experiments on user behavior in cases where systems integrate multiple of such components in a seamless fashion.

2 SYSTEM OVERVIEW

While there is a growing body of research on interactive RS [see e.g. 4, 5], there have been, to our knowledge, no attempts to extend a standard model-based CF recommender into a fully interactive, user-controlled system. Consequently, our previous research was driven by the idea that latent factor models as derived by conventional MF have more potential than currently exploited in RS research. Primarily known for recommendations that appear very precise in offline evaluations, these models have only seldom been used for other purposes. Exceptions include, for instance, preference elicitation [3], diversification [11] or visualization [8].

In [10], we proposed and evaluated a method that presents users with a dialog asking them to choose between sets of items. These sets are automatically generated from an underlying latent factor model: In each step, items are juxtaposed that represent either low or high values for one of the factors. The result is an artificial user-factor vector $\vec{p}\,'_u$ that may be used together with item-factor vectors $\vec{q}_i$ in the dot product to calculate predictions—without forcing users to rate items, as it is customary for CF active learning approaches.

Next, we proposed *TagMF*, a method for integrating standard MF with additional data [2]. First, under the assumption that only content attributes ${}^{i}\mathbf{H}$ for items are known, we redefined the MF model as follows: $\mathbf{R} \approx \mathbf{P}\mathbf{Q}^{\top} = {}^{u}\mathbf{H}\,{}^{u}\mathbf{\Theta}\,{}^{i}\mathbf{\Theta}^{\top}\,{}^{i}\mathbf{H}^{\top}$, with ${}^{u}\mathbf{\Theta}$ and ${}^{i}\mathbf{\Theta}$ associating attributes with factors. Subsequently, we were able to derive the user-attribute matrix ${}^{u}\mathbf{H}$ as well. Implemented by means of tags for movies as a running example, we in [9] not only showed that content-boosting is actually beneficial for users (which previously had only been observed in offline experiments), but in particular, that this regression-constrained formulation allows bringing more interactivity into model-based CF systems and opening up the "black box" the underlying models usually constitute: As the

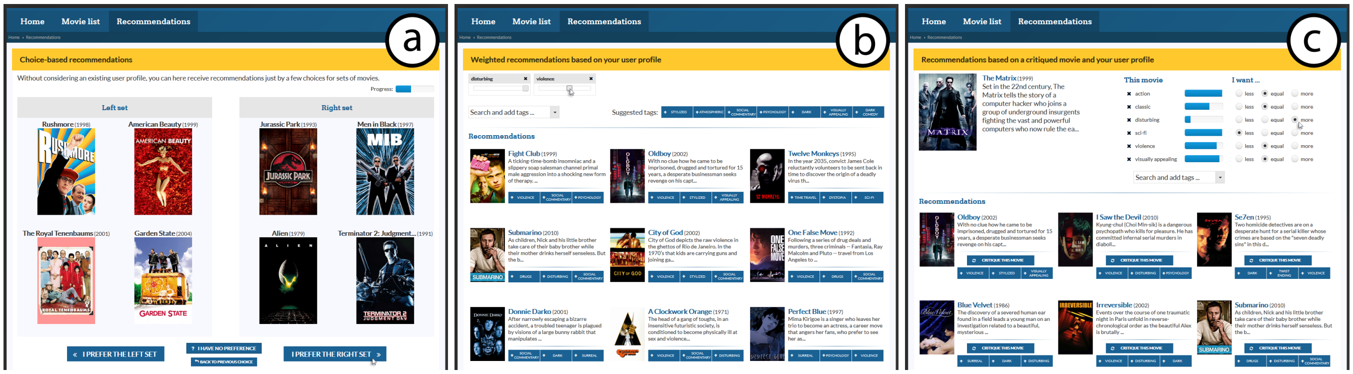

Figure 1: Screenshots of our demonstrator: Users can a) express their preferences through choices, b) manipulate their user-factor vector by selecting and weighting tags, c) apply critiques to request items that represent some aspects less/more strongly.

previously abstract latent factor vectors now comprise well comprehensible information, several promising application possibilities appear that we have described and evaluated in [9], and which form, among others, the core of our novel demonstrator.

3 THE DEMONSTRATOR

For this paper, we implemented all the interactive features we have proposed as separate extensions to MF recommenders in a single web application[1]. Thus, users can create a user profile and receive corresponding recommendations as in conventional CF systems. However, they are provided with additional means to interact with the recommender and to express their preferences, both at cold-start and later in the process. We implemented this demonstrator based on the *TagMF* framework (which we built for this purpose on top of *Apache Mahout* according to the method we have described in [9]) in combination with *MovieLens* datasets and *TMDb* metadata.

Figure 1 shows parts of this application: In (a), the user is confronted with the automatically generated dialog displaying comparisons of item sets representative for factors of the underlying model. After recommendations are generated—based on preferences elicited via this choice-based method, or conventionally by means of a user-factor vector learned off-/online from previously provided ratings—the user can select tags and weight them (b). This way, his or her position in the space spanned by the latent factors is updated, allowing the user to interactively adapt the result set according to the current situation without being required to (re-)rate items. Moreover, when the user clicks on a recommended item to inspect its details, it is possible to critique this recommendation (c): The user can request a new recommendation set containing items that are similar to the currently recommended one, but represent selected dimensions less or more strongly. For methodological details we refer to [9, 10]. All features are connected via several interaction paths, making this and the other features accessible from almost anywhere in the application.

4 CONCLUSIONS & OUTLOOK

In three extensive user studies [9, 10] we have evaluated our proposed interaction and preference elicitation mechanisms. In this context, we have also investigated the impact of additional information in model-based CF on aspects related to user experience in comparison to automated systems exclusively relying on ratings and to interactive tag-based systems [9]. The demonstrator presented in this paper embeds all these features in a single web application using our newly implemented framework, still offering the standard features of CF recommenders. Typical design patterns for RS [1] are taken into account, but adhering to them even more closely is subject of future work. Moreover, we plan to further extend the application, among others, by exploiting information such as user reviews or visual features and by providing more detailed explanations. Beyond that, we especially want to us it as a vehicle for more comprehensive user experiments. The proposed approaches have already been evaluated on their own, but not integrated holistically and based on the latest developments. For instance, we for this paper implemented the choice-based method for the first time on a content-boosted model, making the investigation of possible benefits over the original variant an interesting aspect for future work. In general, we aim at initiating research on user behavior in recommendation scenarios that are as complex as in our demonstrator, i.e. where users can choose from a wide range of options to interact with a system, possibly affected by their situation, their level of expertise, domain knowledge, and other constraints.

[1] http://interactivesystems.info/tagmf

REFERENCES

[1] P. Cremonesi, M. Elahi, and F. Garzotto. 2017. User interface patterns in recommendation-empowered content intensive multimedia applications. *Multimed. Tools Appl.* 76, 4 (2017), 5275–5309.

[2] T. Donkers, B. Loepp, and J. Ziegler. 2016. Tag-enhanced collaborative filtering for increasing transparency and interactive control. In *UMAP '16*. ACM, 169–173.

[3] M. P. Graus and M. C. Willemsen. 2015. Improving the user experience during cold start through choice-based preference elicitation. In *RecSys '15*. ACM, 273–276.

[4] C. He, D. Parra, and K. Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Syst. Appl.* 56, 1 (2016), 9–27.

[5] M. Jugovac and D. Jannach. 2017. Interacting with recommenders – Overview and research directions. *ACM TiiS* 7, 3 (2017), 10:1–10:46.

[6] J. A. Konstan and J. Riedl. 2012. Recommender systems: From algorithms to user experience. *User Model. User-Adap.* 22, 1-2 (2012), 101–123.

[7] Y. Koren, R. M. Bell, and C. Volinsky. 2009. Matrix factorization techniques for recommender systems. *IEEE Comp.* 42, 8 (2009), 30–37.

[8] J. Kunkel, B. Loepp, and J. Ziegler. 2017. A 3D item space visualization for presenting and manipulating user preferences in collaborative filtering. In *IUI '17*. ACM, 3–15.

[9] B. Loepp, T. Donkers, T. Kleemann, and J. Ziegler. 2019. Interactive recommending with tag-enhanced matrix factorization (TagMF). *IJHCS* 121 (2019), 21–41.

[10] B. Loepp, T. Hussein, and J. Ziegler. 2014. Choice-based preference elicitation for collaborative filtering recommender systems. In *CHI '14*. ACM, 3085–3094.

[11] M. C. Willemsen, M. P. Graus, and B. P. Knijnenburg. 2016. Understanding the role of latent feature diversification on choice difficulty and satisfaction. *User Model. User-Adap.* 26, 4 (2016), 347–389.