On User Awareness 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

AWARE Workshop: IUI 2017, March 13-16, 2017, Limassol, Cyprus.

Abstract

In this paper, we discuss several aspects that users are typically not fully aware of when using model-based Collaborative Filtering systems. For instance, the methods prevalently used in conventional recommenders infer abstract models that are opaque to users, making it difficult to understand the learned profile, and consequently, why certain items are recommended. Further, users are not able to keep an overview of the item space, and thus the alternatives that in principle could also be suggested. By summarizing our experiences on exploiting latent factor models for increasing control and transparency, we show that the respective techniques may also contribute to make users more aware of their preferences' representation, the rationale behind the results, and further items of potential interest.

Author Keywords

Recommender Systems; Interactive Recommending; Matrix Factorization; User Interfaces; User Experience.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous.

Introduction

Recommender Systems (RS) that use model-based *Collaborative Filtering* (CF) are known to very efficiently

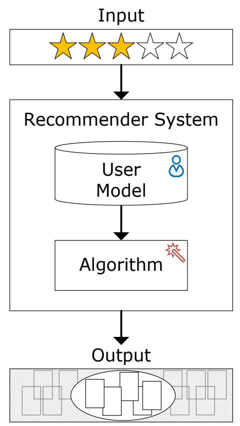

Figure 1: Simplified view on the recommendation process: The user provides feedback, which is used for modeling his or her preferences. The model serves as input to the recommender algorithm. Both representation of preferences and algorithmic internals are usually opaque to the user. Finally, items are presented as recommendations. However, the user is only made aware of a small set of items the system considers most relevant.

generate accurate results. However, high accuracy does not necessarily lead to a commensurate level of user satisfaction. In contrast, user-related aspects such as amount of control users are able to exert over the recommendation process as well as system transparency substantially contribute to acceptance of the results [6]. Still, model-based CF methods such as the popular *Matrix Factorization* (MF) [7] mostly operate as black boxes. Users cannot be aware of how their feedback—either explicitly given or implicitly observed—is taken into account to represent their preferences on side of the system, and consequently to generate recommendations (see Fig. 1). It is therefore often difficult for users to comprehend the impact of their behavior on the underlying model, and the algorithmic rationale for recommending certain items. Further, since results are typically presented in form of ranked lists comprising only items the algorithm considers relevant, it remains unclear how these items relate to the rest of the item space, i.e. items not recommended. Although several attempts have been made to increase inherent diversity of recommendation lists, users' possibilities to actively obtain an overview of the naturally large item space and to become aware of the variety of items that could potentially be recommended, are still very limited. In worst case, this may lead to filter bubble effects.

From our perspective, increasing users' awareness of how their preferences are represented, why particular items are suggested, and which items are among the recommendable alternatives, is hence part of the overall goal of improving RS transparency. In this paper, we will therefore share experiences from prior work on model-based CF, outlining how our attempts to improve user control and system transparency may also contribute to the aforementioned aspects.

Integration of Additional Data

Textual explanations have been widely accepted for supporting transparency of RS [6]. Since model-based CF methods rely on abstract models learned to store user preferences, it is, however, rather difficult to explain their rationale. Due to the statistical nature of latent factor models, this is particularly true when using, e.g., MF [7]. Yet, we argue that *leveraging content-boosted techniques*, i.e. learning an integrated model of latent factors derived from user ratings and additional content information, may raise awareness of *how the system works* and *how the preference profile looks like*.

With *TagMF* [2, 3], we proposed an approach that follows this idea by enhancing MF with relevance data of user-generated tags regarding the items. As shown by others, the additional data positively affects offline accuracy [7]. However, we have also confirmed that it improves perceived recommendation quality, and in particular, transparency [2]. Fig. 2 outlines the findings from the user study (n=46) described in [2], underlining that it appears more clear why items are recommended when preferences are elicited via tags instead of ratings. Further, the interaction possibilities provided in addition to rating items—in CF usually the only means for users to give feedback and to influence the user model—seem to be valuable in terms of perceived control, without increasing the effort [2]. Besides, selecting and weighting tags (Fig. 3), thereby adjusting the usually intransparent user factor vector, may disclose the influence user behavior has on this kind of preference profiles. This, and the ability to revert actions, is important for RS users [5]. Consequently, our approach allows to interactively explore different settings and revoke selected tags anytime, while providing immediate feedback and result updates. Supported by

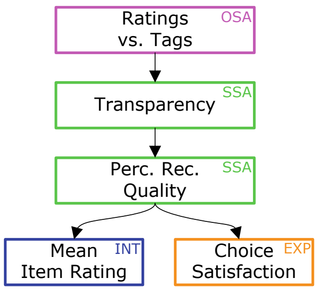

Figure 2: Reduced version of the path model from [2] describing the influence of the objective system aspect (OSA) *preference elicitation via ratings or tags*: The subjective system aspect (SSA) *transparency* is a complete mediator for *perceived rec. quality*, which in turn fully mediates interaction behavior (INT) and user experience (EXP), measured in terms of *mean item rating* and *choice satisfaction*, respectively.

the increased transparency, we argue that users thus become more aware of the system's rationale than when only ratings or implicit feedback is considered.

In the recommendation set, users seem to notice some kind of inner consistency, which is not the case to this extent when results are purely based on ratings. We assume tags to incorporate semantics into the latent factor model that are naturally to understand. In [3], we further investigated the tags' influence. While latent factors are generally considered to describe real-world characteristics [7], our integrated model allows deriving tag-factor relations, thus revealing their actual meaning. As we also derive tag relevances for each user—even if he or she never tagged an item—a user's (former latent) preference profile may now be explicitly described using, e.g., a tag cloud (cf. Fig. 3).

This way offering insights into a latent user model may also help to address privacy concerns. We further assume that the tags shown alongside recommendations (Fig. 3) may be even more expressive when additionally taking latent information into account, thus providing some intuition about the most important factors. These are subjects of future work. Finally, the present results suggest that novelty and diversity do not particularly benefit from integrating tags. Still, as the next section will show, latent factor models may also be exploited to let users discover new alternatives.

Preference Elicitation and Visualizations

Besides explanations and novel preference elicitation methods, also visualizations have been suggested to increase transparency, ranging from diagrams depicting sources in hybrid RS [10] to maps of the item space [4]. Attempts taken in RS research to visualize item space or user profile, however, suffer from several drawbacks [8]. Yet, we argue that solely based on ratings, the *latent factor space* may be exploited to help users *exploring and understanding the item space, its coverage*, and *the representation of their preferences*.

With our choice-based preference elicitation [9], users have to compare sets of sample items that score high or low on several latent factors. This allows users to express their preferences with respect to very diverse items, since each interaction step brings up a different dimension of the factor space. In contrast to other attempts to support users at cold-start, this may upfront increase users' awareness of the variety of options they have in the RS. Consequently, in a user study (n=35), recommended items seemed novel while the approach was found promising in terms of perceived recommendation quality compared to a rating-based attempt.

In [8], we presented a 3D item space visualization where the high-dimensional latent factor space is mapped onto a 2D surface. Preferences estimated for the current user yield the elevation of the landscape (cf. Fig. 4), which in turn can be interactively altered in order to manipulate the underlying profile, i.e. the user factor vector. Shaping hills and valleys allows to express interest for certain areas and to lower the relevance of others. The user study (n=32) we conducted suggests that this helps users understanding the item space, broadening their perspective, and getting an overview [8]. In addition to listing top-n results as in conventional RS, highlighting recommendations inside the landscape supports users in becoming aware of their choice options [8]. Hidden semantics of factors, revealed by *TagMF* by means of tags, here seem to result in a representation users are able to grasp even

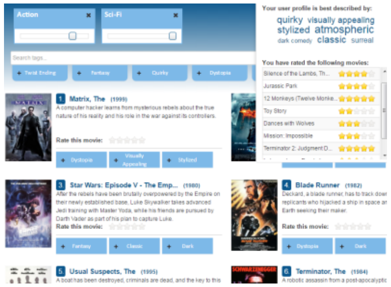

Figure 3: By integrating tags into a latent factor model, users are enabled to understand and manipulate their preference profile expressed implicitly in the (intransparent) latent factor space through explicitly presented textual tags—even if they do not tagged items themselves [2, 3].

Figure 4: Presenting a latent factor space as a landscape (where hills represent areas of the item space the current user prefers and valleys indicate lower relevance) turned out to provide users an adequate overview and to make them aware of alternatives to the recommended items [8].

without integrating additional content data. In search tasks with rather "soft" goals, they were particularly successful, apparently aware of where to look for items e.g. suited for children. Similar to the choice-based approach, this is achieved by showing only a limited number of representative sample items (for areas of the map instead of for single factors). Further, although model-based CF is usually intransparent, reflecting the preference profile as a landscape seems helpful for users to understand how they are represented within the system, and thus, why certain items are recommended.

The aforementioned approaches do not require users to know and rate single items, which is particularly important in unknown domains or for a vague search. However, although empirical results were quite promising [9, 8], search effort may be too high in real-world scenarios. Thus, bringing together *TagMF* with the other approaches might be a valuable opportunity. While *TagMF* could contribute to recommendation quality and transparency in general, it may also be useful to, for instance, further improve exploration by enriching a map visualization with tags labelling the areas.

Conclusions

Overall, it seems important for RS research to face the challenge of making users more aware of, among others, the recommendations at the right time [1], their source in hybrid settings [10], and explanation facilities as well as feedback mechanisms in general, and especially the related consequences for users upon interaction [5]. In this paper, we discussed our prior work on latent factor models in light of user awareness, showing how the respective developments may contribute to enabling users to perceive a number of relevant aspects in a more deliberate fashion.

References

1. Betzalel, N. D., Shapira, B., and Rokach, L. 2015. "Please, not now!": A model for timing recommendations. In *Proc. RecSys '15*. ACM, 297–300.
2. Donkers, T., Loepp, B., and Ziegler, J. 2016. Tag-enhanced collaborative filtering for increasing transparency and interactive control. In *Proc. UMAP '16*. ACM, 169–173.
3. Donkers, T., Loepp, B., and Ziegler, J. 2016. Towards understanding latent factors and user profiles by enhancing matrix factorization with tags. In *Poster Proc. RecSys '16*.
4. Gansner, E., Hu, Y., Kobourov, S., and Volinsky, C. 2009. Putting recommendations on the map: Visualizing clusters and relations. In *Proc. RecSys '09*. ACM, 345–348.
5. Jannach, D., Naveed, S., and Jugovac, M. 2016. User control in recommender systems: Overview and interaction challenges. In *Proc. EC-Web '16*.
6. Konstan, J. A. and Riedl, J. 2012. Recommender systems: From algorithms to user experience. *User Mod. User-Adap.* 22, 1-2, 101–123.
7. Koren, Y., Bell, R., and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *IEEE Computer* 42, 8, 30–37.
8. Kunkel, J., Loepp, B., and Ziegler, J. 2017. A 3D item space visualization for presenting and manipulating user preferences in collaborative filtering. In *Proc. IUI '17*. ACM.
9. Loepp, B., Hussein, T., and Ziegler, J. 2014. Choice-based preference elicitation for collaborative filtering recommender systems. In *Proc. CHI '14*. ACM, 3085–3094.
10. Parra, D., Brusilovsky, P., and Trattner, C. 2014. See what you want to see: Visual user-driven approach for hybrid recommendation. In *Proc. IUI '14*. ACM, 235–240.