
Visualizing Item Spaces to Increase Transparency and Control in Recommender Systems

Johannes Kunkel
University of Duisburg-Essen
Duisburg, Germany
johannes.kunkel@uni-due.de

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

ABSTRACT

Recommender systems (RS) are very common tools designed to help users choose items from a large number of alternatives. While research in RS has been mainly focusing on algorithmic precision, it slowly starts to take user-centric aspects into account as well. In this paper we present two demonstrative applications that target at increasing transparency and control in RS. Both prototypes follow the same method. As a baseline, the entire item space of a domain is visualized using a map-like interface. Inside this depiction users can express their preferences to which the RS reacts with matching recommendations. To change recommendations, users can alter their preferences expressed, which creates a continuous feedback loop between user and RS.

KEYWORDS

Recommender Systems; Interactive Recommending; User Control; Transparency; Filter Bubbles

INTRODUCTION

Recommender systems (RS) are ubiquitous tools for filtering content according to the user's needs. Recommending algorithms have become very mature—being able to determine items that a user will prefer with high accuracy. However, other quality aspects of RS have been taken into consideration only recently.

Most often RS appear as *black boxes* to their users, as it remains opaque why certain items are suggested. This is a well-known issue and can, for instance, be tackled by generating textual explanations for recommendations (as in Amazon's "Customers who bought ... also bought ..."). Even though such

Visualizing Item Spaces to Increase Transparency and Control in Recommender Systems

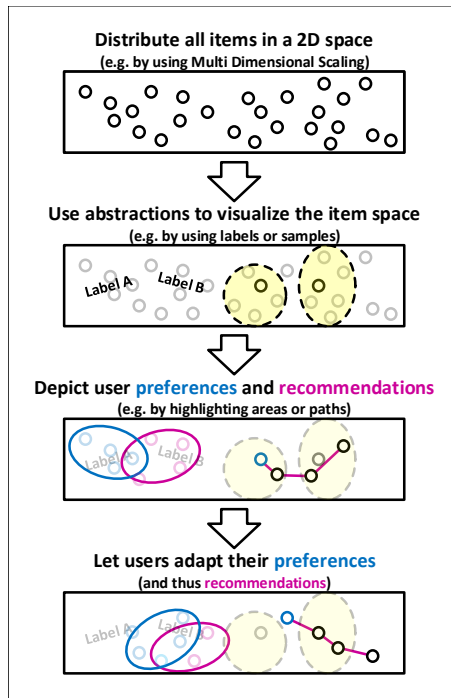


Figure 1: Proposed method for transparent and controllable recommender systems: The item space of a domain is distributed on a 2D plane and visualized using abstractions. Subsequently, user preferences and recommendations can be displayed to users, who can interactively alter preferences and thus recommendations.

approaches can have a positive effect on the perceived transparency, they were criticized for being too shallow [1].

Other approaches that use richer visualizations often target not only a higher transparency, but also the degree of control users can exert [2]. While it was shown that exerting control over a RS can substantially influence user satisfaction [4], many RS lack such means. If present, the only way to adjust underlying preference models (and hence recommendations) is often to re-rate single items, which is cumbersome and lacks efficiency.

Designing more efficient tools to control RS is not trivial, though. Especially when users have only a vague idea of the item domain, it can be difficult for them to enunciate preferences properly. One task of a RS should thus be to educate the user about the underlying item space [3]. Besides enabling adequate expression of preferences, it can also have other benefits when users are aware of the item space. It was, for instance, argued that users who are always presented with a limited subset of items tailored to their specific preferences may get trapped in *filter bubbles*, which may mislead them and even harm society in the long run [6].

In the remainder of this paper, we propose and demonstrate a method for tackling the issues raised by combining state-of-the-art RS with item space visualizations based on maps.

VISUALIZING ITEM SPACES IN RECOMMENDER SYSTEMS USING MAPS

In order to foster transparency and control in RS, we propose a method leveraging the comprehensiveness of geographical maps (see Figure 1) and demonstrate its application using two prototypes: *MovieLandscape* (Figure 2) and *MusicPaths* (Figure 3). Each application is based on a different dataset of items and implements distribution of items, depiction of item space, representation of user preferences and interaction modalities differently.

Distribution of items on a 2D plane: Prior to visualizing the item space, a meaningful distribution of items needs to be found. How this step can be performed depends highly on the background data available. The application *MovieLandscape* (Figure 2) is based on a dataset containing about 10 000 movies and 20 million numerical user ratings. Utilizing these ratings, similarities between all items are calculated and translated to x/y coordinates using Multi Dimensional Scaling [5]. The resulting distribution places items that are rated similarly near each other and items that are rated very differently on opposite sides of the map, thus leveraging the hidden semantics in users' rating behavior. *MusicPaths* (Figure 3), on the other hand, makes heavy use of additional information regarding the content of items. The background dataset contains about 630 000 music songs tagged with 2 mood tags each. For distribution, all existent moods are laid out on a plane and songs are placed according to their associated mood tags. The moods are arranged in such a way that dimensions correspond to a continuum between *calm* and *energetic* on the x axis, and *dark* and *positive* on the y axis.

Visualizing Item Spaces to Increase Transparency and Control in Recommender Systems

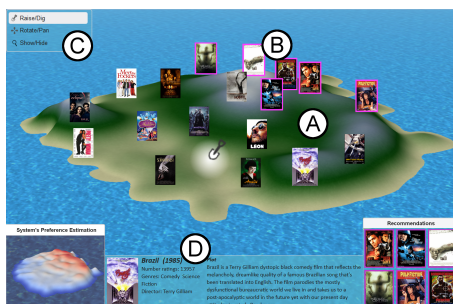


Figure 2: The *MovieLandscape* [5] application with representation of item space (A); user's preferences indicated by surface elevation; recommended items (B) shown with magenta-colored margins; a palette (C) with tools enabling interaction and details of a currently selected movie (D).



Figure 3: The *MusicPaths* application with the item space being represented as an island of moods (A); a current playlist of songs depicted as path (B); an interactive compass for steering the direction the playlist develops into (C); and details of the currently played song (D).

Abstraction of item space: Since the entire item domain will most likely be too large to be directly comprehensible (item spaces typically contain thousands of items), an abstraction must be chosen that adequately conveys the arrangement of items. In *MusicPaths* the map is divided into areas that are labeled according to the underlying distribution of mood tags (Figure 3, A), whereas in *MovieLandscape* a clustering is performed in order to identify groups of items that are similar to each other. For each of these clusters, one representative sample item is chosen, which is presented by its movie poster (Figure 2, A).

Depiction of user preferences and recommendations: The domain and structure of background data together with the chosen visualization of the corresponding item space set the baseline for depicting preferences and recommendations. For instance, music is typically consumed as a sequence, which should be considered when depicting song recommendations. Consequently, in *MusicPaths* the active user defines a point on the map and a direction, on which the system reacts with a path of music songs starting from the position given and heading in the specified direction (Figure 3, B). This setting also aligns well with the used item space visualization based on moods. Starting from a current mood, a path denotes towards which mood the playlist develops. The interaction with *MusicPaths* thus targets a roaming behavior where the user wanders over the map of music while listening to different songs, potentially changing the mood he or she is currently in. Such a scenario is not as promising for situations where only one item is searched. Hence, *MovieLandscape* shows preferences as areas (landscape elevations) and recommendations as additional samples inside these parts (Figure 2, B).

Interactively adapt preferences and recommendations: One of the most important aspects we wanted to address is the control users can exert over the RS. Again, the form of exerting control depends on how item space and preferences are expressed to the user. In order to stick as close as possible to the metaphorical language of roaming a map with music, in *MusicPaths* users can adjust the direction in which they wish to wander over the map by using an interactive compass (Figure 3, C). By clicking on one of the eight compass directions, users are able to change the general direction of the playlist generation. In *MovieLandscape*, on the other hand, users can express preferences by reshaping the 3D landscape so that hills indicate areas of high preference and valleys indicate areas with items the users do not like. In both cases, user interaction triggers an immediate recalculation and depiction of recommendations.

In order to test our prototypes and the general method behind them, we conducted user studies for *MovieLandscape* ($N=32$) and *MusicPaths* ($N=30$). Some results can be found in Table 1 and further in a previous publication [5]. Altogether, results show that study participants liked our approaches and—most importantly—felt in control over their recommendations (as expressed by the construct *interaction adequacy* [7] in Table 1). However, some aspects undoubtedly show potential for improvement. While *MovieLandscape* scores rather good in terms of perceived transparency, control and overall satisfaction, results for *MusicPaths* are consistently lower. We ascribe this to the inferior means of scrutinizing the

Table 1: Descriptive results for the applications *MovieLandscape* (ML) and *MusicPaths* (MP), respectively. All constructs belong to the *ResQue* inventory [7] and were assessed using 5-point Likert scales.

	ML		MP	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Transparency	3.91	1.09	3.30	0.95
Interaction adequacy	3.47	0.88	3.01	0.83
Overall satisfaction	3.94	0.76	3.27	1.05

item space in *MovieLandscape* compared to *MusicPaths*. This is underlined by textual comments of some participants, who wished further interaction tools to explore the music domain (e.g. tools for zooming into the map and revealing songs in specific areas). Others also disagreed with the assignment of moods to songs and wished to replace them. This may have lowered perceived transparency and, as a consequence, made the RS harder to control. Altogether, results may also indicate that music domains are in general harder to depict in a comprehensible manner.

CONCLUSIONS AND FUTURE WORK

We propose a novel method for presenting users with visualizations of item spaces, their preferences and recommendations, utilizing a map metaphor. Two prototypes demonstrate applicability of the approach in the domains of movies and music. First user tests indicate a positive user response, but reveal potential for improvements too. A crucial point is the comprehensibility of the map layout. An item space map that appears obscure to its users is very likely to narrow down the perceived transparency, controllability and general usability of the entire application. Future developments will thus aim at revealing factors that make such map-like visualizations of item spaces more comprehensible. We also plan to systematically analyze dependencies between stages, for instance, to reveal general relations between structure of background data and depiction of item space.

REFERENCES

- [1] Shlomo Berkovsky, Ronnie Taib, and Dan Conway. 2017. How to Recommend?: User Trust Factors in Movie Recommender Systems. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces (IUI '17)*. ACM, New York, NY, USA, 287–300. <https://doi.org/10.1145/3025171.3025209>
- [2] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. *Expert Systems with Applications* 56 (2016), 9–27. <https://doi.org/10.1016/j.eswa.2016.02.013>
- [3] Dietmar Jannach and Gediminas Adomavicius. 2016. Recommendations with a Purpose. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. ACM, New York, NY, USA, 7–10. <https://doi.org/10.1145/2959100.2959186>
- [4] Michael Jugovac and Dietmar Jannach. 2017. Interacting with Recommenders—Overview and Research Directions. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 7, 3 (2017), 10:1–10:46. <https://doi.org/10.1145/3001837>
- [5] Johannes Kunkel, Benedikt Loepp, and Jürgen Ziegler. 2017. A 3D Item Space Visualization for Presenting and Manipulating User Preferences in Collaborative Filtering. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces (IUI '17)*. ACM, New York, NY, USA, 3–15. <https://doi.org/10.1145/3025171.3025189>
- [6] Eli Pariser. 2011. *The filter bubble: What the Internet is hiding from you*. Penguin UK.
- [7] Pearl Pu, Li Chen, and Rong Hu. 2011. A User-centric Evaluation Framework for Recommender Systems. In *Proceedings of the Fifth ACM Conference on Recommender Systems (RecSys '11)*. ACM, New York, NY, USA, 157–164.