

Towards a Design Space for Personalizing the Presentation of Recommendations

Catalin-Mihai Barbu and Jürgen Ziegler

University of Duisburg-Essen
Duisburg, Germany

`catalin.barbu@uni-due.de`, `juergen.ziegler@uni-due.de`

Abstract. Although personalization plays a major role in the development of recommender systems, the presentation of recommendations—and especially the way in which it can be adapted to suit the user’s needs—has received relatively little attention from the research community. We introduce a design space for personalizing the presentation of recommendations and propose several dimensions that should be a part of it. Moreover, we present our initial insights about possible interactive mechanisms as well as potential evaluation criteria. Our goal is to provide a systematic way of designing personalized recommendation content, which should prove beneficial for other researchers working on this topic. In the longer term, we are interested to investigate whether such personalized presentation implementations influence the perceived trustworthiness of the recommendations.

Keywords: Recommender systems; personalization; design space; interactive control

1 Introduction & Motivation

Personalization is an important aspect of recommender systems (RS). It allows websites and other Internet services to cater to individual tastes, interests, and preferences. For many years, objective accuracy was considered one of the most important criteria for ranking RS [11]. Consequently, the use of personalization was mostly focused on improving the algorithms and models used to generate result sets. However, recommendations are only as good as users perceive them to be. More recently, some researchers have begun to argue that subjective accuracy is equally, if not more, important than objective accuracy and may play a larger role in determining user satisfaction [3, 11]. Perceived accuracy has been shown to be influenced positively by user-related aspects such as control, trust, and transparency [3, 13]. Personalization is already one of the methods used to help users understand why a recommendation is suitable for them. Previous research has investigated its positive influence on user experience [10, 17]. Combining personalization techniques with novel approaches from the field of interactive RS could therefore lead to additional insights into how user satisfaction can be increased even further.

A relatively unexplored topic in the field of RS is the personalization of the presentation of recommended items. Once user preferences have been elicited (either implicitly or explicitly), this information can be used not only to suggest personalized predictions, but also to customize the way in which they are presented to the user. Adapting the presentation to fit the user’s needs has the potential to open novel interaction possibilities for users and might provide useful insights into the way in which people interact with RS. Against this background, exploring the design space for the personalization of recommendations is a useful research endeavor and an important step towards the implementation of a prototype. The goal of this paper is to introduce a design space for personalizing the presentation of recommendations and to present the dimensions that comprise it.

The remainder of the paper is structured as follows: We discuss related work in Section 2, before proceeding to present the design space in Section 3. We subsequently introduce some preliminary interactive mechanisms and evaluation criteria. Finally, we discuss possible limitations and directions for future research in Section 4.

2 Related Work

Personalization is well-studied in the field of RS. Some of the main research foci include deciding, for a given recommendation, what information to present, when to present it, how much of it to present, and in what way. For instance, different information modalities (such as various types of result lists or combinations of text and images) have been compared to observe their effect on the persuasiveness of recommendations and on the users’ satisfaction [12]. Prior work has also investigated models for context-aware RS that can predict the best time to show recommendations [5]. Other researchers have determined the number of items in a result set that maximizes choice satisfaction without increasing choice difficulty [1].

Many existing approaches to personalizing the presentation of recommendations rely on explanations [13, 16, 19]. “Common sense” approaches, which use rules to determine what items to recommend and how to personalize the presentation have also been developed [6]. Novel approaches for visualizing recommendations have been proposed, such as those implemented in *TasteWeights* [2] and *TalkExplorer* [18]. These interactive approaches afford a certain degree of control over the recommendation process to elicit feedback and preferences as well as to increase transparency. The effects of personalization, especially with respect to the use of explanations, have been investigated in several prior works (see, e.g., [15] and [17]).

Previous research into design spaces for adaptive user interfaces highlighted the importance of control over the adaptation algorithm and the importance of adequate measures for user evaluation [8]. This research focused on generic user interface control structures (e.g., menus) and did not cover information systems such as RS. To the best of our knowledge, no prior work has so far

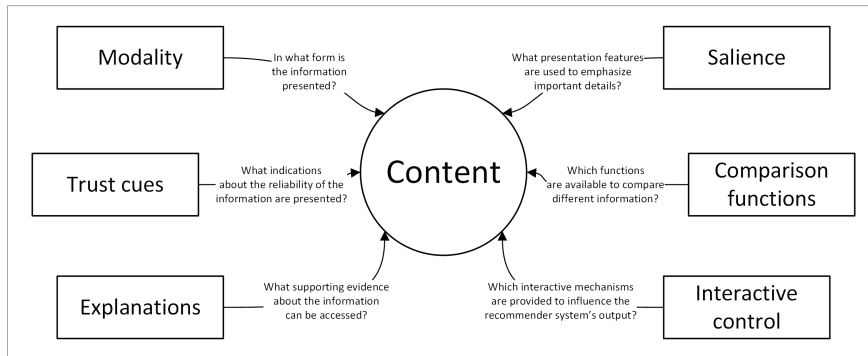


Fig. 1. Overview of design space.

focused explicitly on exploring the design space of personalized presentation of recommendations.

3 Analysis of Design Space

We identify the following dimensions that comprise the design space (Fig. 1): modality, saliency, comparison functions, interactive control, explanations, and trust cues. Each of these is explained in further detail below.

The design space is meant to be applicable to numerous domains in which RS are commonly used. To facilitate understanding of the various dimensions, throughout this section we limit ourselves to using examples from the hotel booking domain. Hotel recommendations are interesting for several reasons. First, there is a higher risk associated with such choices—in comparison with movie recommendations, for example. Risk arises, on the one hand, from the fact that staying in a hotel typically costs a considerable amount of money. On the other hand, there is also the risk associated with the effects of a wrong recommendation on the user’s wellbeing. Second, the items in question have a reasonable set of attributes that should be considered. These can be classified into hotel features (e.g., location, price), room characteristics (e.g., bed size, number of electrical sockets), and services (e.g., complimentary breakfast, free Wi-Fi). Third, there is a large body of user-generated content, in the form of reviews, photos, tags, and ratings, that can be leveraged in the presentation.

3.1 Design Space Dimensions

Modality refers to the form in which the content of the recommendation is conveyed to the user. Information can be presented using text, graphical symbols, audiovisual means, or combinations thereof [4]. Finding the most appropriate modality for each type of content (for example, description, ratings, pricing information, user reviews etc.) is an important aspect of personalization [12]. Some

users prefer to read an exhaustive description of the hotel to decide whether it matches their requirements; others like viewing photos of the property. Furthermore, some modalities may not be suitable for users with visual or auditory impairments. Changing information modalities may also require that the system adopt a different recommendation paradigm.

Saliency denotes the range of presentation features that are used to draw users' attention. Particularly relevant information, such as attributes in which a person is interested, should be emphasized. Conversely, less important features might be shown in a subtler manner or even hidden altogether (e.g., business services for vacationers). Standard presentation layouts, such as category-value tables, can become difficult to parse if they exceed a certain number of rows. Similarly, altering the size and color of text or using animations to highlight important aspects can lead to information overload if used excessively. Instead, a RS might re-order the list of attributes such that those that the user considers most important are displayed at the top [9]. Relevant additional information can also be shown directly. An example would be displaying the opening times of the local gym to users who have expressed interest in fitness (as opposed to simply listing "fitness center" as a hotel amenity).

Comparison functions help users evaluate item attributes and values across different recommendations [9]. For example, consider people who enjoy spacious accommodations. When browsing hotel recommendations, they would, presumably, look specifically for details about the size of the rooms. The same information might be presented differently by various vendors: as an area (e.g., "14 m^2 "); as the product of individual dimensions (e.g., "4x3.5 m ."); using different units (e.g., "150 sq ft "); in relative terms (e.g., "standard size"); using a blueprint on which the layout and dimensions of the room are depicted. In other cases, such details might be missing altogether. A low comparability has a detrimental effect on the user's decision making processes as well as on her trust in the generated results. The RS should therefore adapt the presentation such that attribute values are normalized to facilitate comparison.

Interactive control comprises the mechanisms through which a user influences the output of a RS. The complexity of the underlying algorithms that are used to generate recommendations has increased a great deal over the years. For this reason, many users associate modern RS with "black boxes" [13]. The lack of transparency and limited options for controlling the output are frequently cited as reasons for the users' lack of trust in the recommended items [13]. Various approaches for increasing user control have been proposed, ranging from novel ways to elicit preferences [7] to innovative frameworks for enhancing decision support [9]. A straightforward example would be a hotel RS that allows users to modify the relative weights of hotel attributes per their own preferences.

Explanations allow users to discover supporting evidence for a presented attribute or claim and are one of the more common methods for increasing transparency in RS [14, 16, 19]. A hotel description claiming that the establishment is close to the city center may be misleading. It might measure only the distance to the edge of the central district (rather than the geographical center),

or simply provide a “straight-line” distance that is of little help in practice. If location is an important aspect for the user, the RS might, for example, display the average walking time based on information extracted from user reviews or from local transportation websites. More interactive approaches could leverage GPS data to display a map that allows her to calculate the travel time between the hotel and various landmarks, perhaps even using various means of transportation. Providing sufficient evidence is important for both objective as well as subjective information. The former might be wrong or incomplete, whereas the latter might need to be put into the proper context. Explanations could also help clarify why a certain piece of information is presented—as well as why it is presented in a certain way. To achieve this, the RS should be able to represent the user’s personalization profile in a meaningful way.

Trust cues are interface elements that allow the user to determine the reliability of the presented information [16]. Item descriptions should be complemented, to the extent possible, by objective measurements. The credibility of user-generated content, such as user reviews, should also be evaluated. When personalizing the presentation of a recommendation, a RS might show supporting evidence contributed by trustworthy reviewers. This means ensuring, on the one hand, that a review is not fake, and on the other hand, that the reviewer has sufficient expertise. It is, however, equally important to recognize that the system has limited knowledge of its users’ (personalization) preferences. Hence, the RS should provide adequate trust cues to make the user aware of this inherent uncertainty. In other words, a person who is considering a recommendation should understand how trustworthy each part of the recommendation is. This ensures that the user’s perceived trustworthiness of the RS remains in sync with the system’s actual trustworthiness [16]. As an example, consider the case where multiple reviewers have complained about the stiffness of the bed in a hotel that otherwise appears to be a good match for a prospective traveler. The reliability of this piece of information depends on how long ago the reviews were written, on the proportion of guests who made similar comments, as well as on their breadth of travel experience. All this should be considered and presented to the user in a transparent manner.

3.2 Interactive Mechanisms

A promising approach for personalizing the presentation of recommendations is to employ interactive mechanisms that support the user’s decision making process [9]. Instead of simply ordering reviews by date, a RS might preselect reviews from people who have commented on issues that match the user’s interests. Initially, only the most relevant comments would be shown, though the user would be afforded the option to expand each review fully. Going a step further, a RS might offer “personalized summaries” containing relevant attributes, aggregate ratings, review snippets, as well as relevant photos or maps.

The content of recommendations is typically organized into sections, such as general description, listing of attributes, ratings, tags, reviews or comments, and photos. It is reasonable to expect that users’ preferences extend also to the

order in which these sections are presented. For example, one might consider user reviews more relevant than the owner-supplied description of the hotel. Hence, the RS could allow users to customize the various content sections to their liking. Furthermore, the system might also attempt to match users' expectations based on available context information.

Further interactive mechanisms could be developed to facilitate users' control over their own personalization profiles. Ideally, the system should not only provide the means for users to edit their profiles, but also to preview the effect that a prospective change would have on the presentation of the recommendations.

3.3 Evaluation Criteria

Based on the dimensions presented above, it seems possible to devise methods for evaluating RS with respect to how strongly the different dimensions are perceived by users. We believe that the main criterion for evaluating user interfaces that implement the design space is their *suitability with respect to the user's informational need*. This may depend on several factors, such as the consequences of choosing wrongly (e.g., in terms of monetary costs or the user's wellbeing), the required level of detail (i.e. how accurate does the information need to be), and user characteristics.

4 Discussion and Future Work

Personalizing the presentation of recommended items using interactive mechanisms can lead to increased transparency and control over the recommendation process. Since both aspects are central to the issue of trust, this additional kind of personalization might increase the perceived trustworthiness of the recommendations. However, this will need to be investigated empirically.

A limiting factor for the design of such interactive and personalized presentations is the quality of the user data, such as elicited preferences, that is available to the RS. At the same time, many of the existing user models are not optimized sufficiently to support this level of customization. Therefore, one of our planned directions for future research is to investigate how the user models commonly used in RS can be expanded to afford personalized presentations of recommendations.

Although we decided to focus in this paper only on the presentation of individual recommendations, personalized presentations also make sense for the item sets (i.e. even before the user has selected a recommendation for closer scrutiny). We believe that the design space can be employed successfully in this situation, though a more thorough examination of this case is required.

Finally, we plan to validate the design dimensions presented in this paper and to develop a prototype implementation of the design space on top of an existing platform for hotel recommendations.

Acknowledgments

This work is supported by the German Research Foundation (DFG) under grant No. GRK 2167, Research Training Group “User-Centred Social Media”.

References

1. D. Bollen, B. P. Knijnenburg, M. C. Willemsen, and M. Graus. Understanding choice overload in recommender systems. In *Proceedings of the Fourth ACM Conference on Recommender Systems*, RecSys '10, pages 63–70, New York, NY, USA, 2010. ACM.
2. S. Bostandjiev, J. O'Donovan, and T. Höllerer. Tasteweights: A visual interactive hybrid recommender system. In *Proceedings of the Sixth ACM Conference on Recommender Systems*, RecSys '12, pages 35–42, New York, NY, USA, 2012. ACM.
3. Y. Chen. Interface and interaction design for group and social recommender systems. In *Proceedings of the Fifth ACM Conference on Recommender Systems*, RecSys '11, pages 363–366, New York, NY, USA, 2011. ACM.
4. P. Cremonesi, M. Elahi, and F. Garzotto. User interface patterns in recommendation-empowered content intensive multimedia applications. *Multimedia Tools Appl.*, 76(4):5275–5309, Feb. 2017.
5. N. Dali Betzalel, B. Shapira, and L. Rokach. ”please, not now!”: A model for timing recommendations. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys '15, pages 297–300, New York, NY, USA, 2015. ACM.
6. B. De Carolis, I. Mazzotta, N. Novielli, and V. Silvestri. Using common sense in providing personalized recommendations in the tourism domain. In *In Proceedings of the Workshop on Context-Aware Recommender Systems*, 2009.
7. T. Donkers, B. Loepp, and J. Ziegler. Tag-enhanced collaborative filtering for increasing transparency and interactive control. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization*, UMAP '16, pages 169–173, New York, NY, USA, 2016. ACM.
8. L. Findlater and K. Z. Gajos. Design space and evaluation challenges of adaptive graphical user interfaces. *AI Magazine*, 30(4):68–73, 2009.
9. A. Jameson, M. C. Willemsen, A. Felfernig, M. de Gemmis, P. Lops, G. Semeraro, and L. Chen. Human decision making and recommender systems. In F. Ricci, L. Rokach, and B. Shapira, editors, *Recommender Systems Handbook*, pages 611–648. Springer, 2015.
10. B. P. Knijnenburg, M. C. Willemsen, Z. Gantner, H. Soncu, and C. Newell. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction*, 22(4-5):441–504, Oct. 2012.
11. J. A. Konstan and J. Riedl. Recommender systems: From algorithms to user experience. *User Modeling and User-Adapted Interaction*, 22(1-2):101–123, Apr. 2012.
12. T. Nanou, G. Lekakos, and K. Fouskas. The effects of recommendations' presentation on persuasion and satisfaction in a movie recommender system. *Multimedia Syst.*, 16(4-5):219–230, Aug. 2010.
13. J. O'Donovan and B. Smyth. Trust in recommender systems. In *Proceedings of the 10th International Conference on Intelligent User Interfaces*, IUI '05, pages 167–174, New York, NY, USA, 2005. ACM.

14. P. Pu and L. Chen. Trust-inspiring explanation interfaces for recommender systems. *Know.-Based Syst.*, 20(6):542–556, Aug. 2007.
15. P. Pu, L. Chen, and R. Hu. Evaluating recommender systems from the user’s perspective: Survey of the state of the art. *User Modeling and User-Adapted Interaction*, 22(4-5):317–355, Oct. 2012.
16. D. Sacha, H. Senaratne, B. C. Kwon, G. P. Ellis, and D. A. Keim. The role of uncertainty, awareness, and trust in visual analytics. *IEEE Trans. Vis. Comput. Graph.*, 22(1):240–249, 2016.
17. N. Tintarev and J. Masthoff. Designing and evaluating explanations for recommender systems. In F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, editors, *Recommender Systems Handbook*, pages 479–510. Springer, 2011.
18. K. Verbert, D. Parra, P. Brusilovsky, and E. Duval. Visualizing recommendations to support exploration, transparency and controllability. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces*, IUI ’13, pages 351–362, New York, NY, USA, 2013. ACM.
19. J. Vig, S. Sen, and J. Riedl. Tagsplanations: Explaining recommendations using tags. In *Proceedings of the 14th International Conference on Intelligent User Interfaces*, IUI ’09, pages 47–56, New York, NY, USA, 2009. ACM.