

User Model Dimensions for Personalizing the Presentation of Recommendations

Catalin-Mihai Barbu
University of Duisburg-Essen
Duisburg, Germany
catalin.barbu@uni-due.de

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

ABSTRACT

Personalization in recommender systems has typically been applied to the underlying algorithms and to the predicted result sets. Meanwhile, the presentation of individual recommendations—specifically, the various ways in which it can be adapted to suit the user’s needs in a more effective manner—has received relatively little attention by comparison. 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 recommender system. At the same time, many of the existing user models are not optimized sufficiently for this specific type of personalization. We present the results of an exploratory survey about users’ choices regarding the presentation of hotel recommendations. Based on our analysis, we propose several novel dimensions to the conventional user models exploited by recommender systems. We argue that augmenting user profiles with this range of information would facilitate the development of more interactive mechanisms for personalizing the presentation of recommendations. This, in turn, could lead to increased transparency and control over the recommendation process.

Author Keywords

Recommender systems; personalization; user profile

ACM Classification Keywords

H.5.2 [Information Interfaces and Presentation]: User Interfaces—*evaluation/methodology, graphical user interfaces (GUI), user-centered.*

INTRODUCTION & MOTIVATION

Personalization is an important and well-studied topic in recommender systems (RS). A non-personalized RS [12] will show the same set of recommendations to everyone (e.g., the ten most popular products on an e-commerce site). In contrast, personalized systems allow websites and other Internet services to cater to individual tastes, interests, and preferences. This is achieved, in part, by exploiting user information collected during the interaction with the RS. Previous research has noted the positive effect of personalization on enhancing user experience [7]. Personalization also has the potential to increase knowledge

about the domain in which the RS is used, support decision-making processes (e.g., by presenting information that would not otherwise be known to the user), and might even play a role in increasing people’s trust in the recommended items as well as in the RS itself.

Personalizing the presentation of recommended items is still a relatively open topic in the field of RS. Once user preferences have been elicited (either implicitly or explicitly), this information can be used not only to offer personalized predictions [6], but also to customize the way in which these predictions are presented to the user. Adapting the presentation to fit the consumer’s needs has the potential to open novel interaction possibilities [8]; and it might provide useful insights into the various ways in which people interact with such systems. The goal of this paper is to introduce several novel dimensions to the conventional user models exploited by RS. We believe this could facilitate the development of more interactive mechanisms for personalizing the presentation of recommendations.

The remainder of the paper is structured as follows: We discuss related work on personalization and user models in Section 2, before proceeding to present the design and results of the exploratory study in Section 3. We introduce our proposed user model dimensions for personalizing the presentation of recommendations in Section 4. Finally, we draw future research directions in Section 5.

RELATED WORK

Some of the main research foci of personalization 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 [9]. Prior work has also investigated models for context-aware RS that can predict the best time to show recommendations [1]. Other researchers have determined the number of items in a result set that maximizes choice satisfaction without increasing choice difficulty [2].

Many existing approaches to personalizing the presentation of recommendations rely on explanations [13,10,15]. “Common sense” approaches, which use rules to determine what items to recommend and how to personalize the

presentation have also been developed [4]. Novel approaches for visualizing recommendations have been proposed, such as those implemented in *TasteWeights* [3] and *TalkExplorer* [14]. 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 [11,13].

EXPLORATORY STUDY

We conducted an exploratory online study to investigate participants' choices about hotel booking. In selecting the domain, we considered three aspects: 1) The choice should carry a substantial amount of risk for the user; 2) the items should have a reasonable set of attributes that need to be considered; and 3) there should be a large body of user-generated content available, in the form of reviews, photos, tags, and ratings, that can be leveraged for the presentation. Because of the first criterion, we decided to focus on hotel recommending—as opposed to the more common domain of movie recommendations.

Study Design

We theorize that the way in which people make decisions about hotel booking, their trust in social media, and their travel habits influence the information they want to see in a recommendation (i.e. the type of personalization they expect). Our aim for this study was to investigate whether the travel scenario influences users' decision-making processes in ways that can be used to personalize the presentation of hotel recommendations.

The survey was organized in six parts. The first four sections elicited answers regarding our participants' demographics, trust in social media, experience with hotel booking portals, and travel behavior. A filter question was used to assign each participant to one of five travel scenarios: city break / short vacation (1-2 nights), short business trip (1-2 nights), long vacation (3+ nights), long business trip (3+ nights), or family vacation (with children).

In each scenario, users were presented with an identical mockup of a hotel recommendation (Figure 1). First, participants were asked to rank each section of the mockup—overall rating, price, general description of the hotel, photos, a map showing the hotel's location within the city, nearby transportation options, hotel and room amenities, and reviews from users—depending on how important they considered the information in that section to be. Second, they had to select up to 7 topics about which they would like to receive more information when looking at recommendations (e.g., proximity to public transport, room sizes and layouts, or fitness center equipment).

Finally, participants were asked 12 questions designed to determine their typical decision-making behavior during hotel booking. This part was modelled based on the Rational-Experiential Inventory [5], which is designed to measure participants' need for cognition and faith in

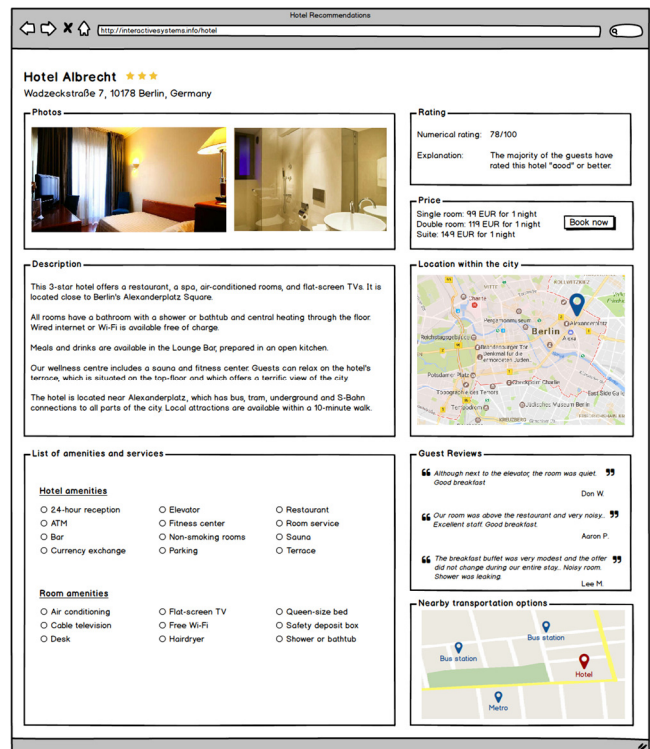


Figure 1: Hotel mockup used in the exploratory study.

intuition, respectively. The questions addressed six underlying factors: a) perceived effort required to complete a hotel booking task; b) economic considerations; c) clearness of mental goal; d) self-efficacy (i.e. trust in one's own choices); e) influenceability; and f) engagement. Each factor was tested through two questions: one with a high and one with a low factor loading, respectively.

Study Results

The survey was published online in January 2017 and ran for one month. A total of 159 participants (82 female; median age in the interval 25-34 years) completed the survey fully. Of the respondents, 123 (77.36%) were employed and 24 (15.09%) were students. Furthermore, 139 (87.42%) had completed at least a Bachelor education. As monetary incentive, all complete responses entered a raffle for one of four Amazon gift vouchers, each worth 25 EUR.

Apart from “family vacation”, all other scenarios were selected by enough participants to allow for meaningful statistical measurements. Most participants (51%) rated their trust in online reviews as high or very high on a 5-point Likert scale ($M=3.53$, $SD=0.71$). These findings were similar across all scenarios. After data analysis (ANOVAs with Fisher's LSD), we noticed a significant difference ($p < 0.05$) when comparing the business scenarios: Over 65% of participants whose typical travel scenario was “long business trip” reported a high or very high trust in online reviews, compared to only 48% in the “short business trip”. The availability of reviews was rated as very or extremely useful by 78% of participants ($M=3.96$, $SD=0.75$). Similarly, photos were considered very or extremely useful

by 82% of respondents ($M=4.17$, $SD=0.82$). In both cases, we observed no significant differences between travel scenarios. We also investigated which characteristics make reviews helpful. An overwhelming majority (92%) stated that useful reviews mention both positive and negative aspects. Furthermore, reviews should be sufficiently detailed (57%), credible (53%), and should give the impression that the reviewer is knowledgeable about the subject (52%).

Certain patterns emerged with respect to users' typical decision-making behavior during hotel booking. First, booking a hotel for vacation is considered more challenging than for business travel—especially for longer stays ($p < 0.05$). Second, people who typically go on longer vacations need more time to decide which recommendation to follow when prices are higher than they are used to. The difference was significant ($p < 0.05$) when compared to the answers from the “long business trip” scenario. Third, participants tend to revisit recommendations to ensure they do not miss important information. A significant difference ($p < 0.01$) was observed when comparing the scenarios “short vacation” and “short business”. These results suggest that the travel scenario can be a factor for personalizing the presentation of recommendations. However, its influence may be lower than predicted (Figure 2).

Contrary to our expectations, we observed almost no significant differences in terms of the importance of the mockup sections for the different scenarios. The sole exception ($p < 0.01$) was “general hotel description”, which proved particularly unimportant for respondents in the “long business” scenario. Similarly, the list of topics about which participants stated they would like to see more information when browsing recommendations did not exhibit significant differences across scenarios.

USER MODEL DIMENSIONS

The exploratory survey provides some initial insights about how user models could be enhanced to facilitate the personalization of the presentation of recommendations. We describe each proposed dimension—information need, personal risk profile, engagement, speed of decision, and trust in social media—separately.

Personalization requires a good understanding of the *user's informational need*, which may, in turn, depend on several factors. If the consequences of choosing wrongly are high (e.g., in terms of monetary costs or the user's wellbeing), the informational need for the user will likely also be greater. This matches our finding that people spend more time looking at hotel recommendations when the prices are higher. Another factor is the required level of detail (i.e. how accurate the information needs to be). This would allow a RS to decide, for instance, whether to show a brief comment about a hotel's location or a longer description that includes nearby points of interest and transportation options. Finally, user characteristics, such as previous experience with hotel booking, degree of trust in the system, or self-efficacy might also play a role in defining the information need.




Figure 2: Results of users' decision-making behavior during hotel booking. Error bars denote the 95% confidence interval.

As we have hinted previously, choosing a suitable hotel from amongst several recommendations is a decision problem that involves a significant amount of risk—regardless of whether the person is planning a business or a leisure trip. Other domains have similar risks associated with such choices. It is therefore necessary to also consider the user's *personal risk profile*. This comprises attributes that could have a detrimental effect on the user's wellbeing if they were to occur in a recommendation that the user ends up following. For example, a hotel in which the beds have particularly stiff mattresses might be problematic for a person who suffers from chronic back pain. A particularly interesting situation arises when such a hotel would otherwise be a very good match for the user. An interactive RS could try to preempt the possibility of a bad choice by compiling a list of complaints based on reviews written by previous guests.

The user's *engagement* with the RS may also be modeled. This refers to the amount of time and effort that a person is willing to spend looking for recommendations. Based on the results of our exploratory study, it seems likely that people browsing hotel recommendations for an upcoming vacation may be more willing to spend time finding the best option. The difference might be caused by the fact that stricter rules typically apply for business travel. For example, the price range may be well-defined and constraints regarding the location might apply. Users might also not have very much time at their disposal for making a choice, thereby opting for a satisficing strategy. The RS might exploit this knowledge to decide which parts of a recommendation to make more salient and which modalities are best suited for presenting certain information about the hotel.

A similar user model dimension is the *speed of decision*, i.e. how quickly the user decides which recommendation to follow. The RS might adapt the presentation of specifically

to support people who find it more difficult to reach a decision (for example, novice users, or those who travel seldom). Representative ways to achieve this could be to ensure that the attributes of the individual recommendations are easy to compare (e.g., by transforming and normalizing units, or always listing attributes in the same order), that the most important characteristics of the hotel are summarized to facilitate quick consumption, or that enough trust cues are present to allow the user to verify the information.

The user's *trust in social media* could also be used to adapt the output of RS. In our exploratory study, most participants expressed their confidence in online reviews, provided they exhibit several characteristics, as mentioned in the previous section. Two aspects are worth mentioning here: First, if the user's trust in social media is low, the RS might allocate less space to user reviews and only show the most credible ones (e.g., those written by experienced reviewers). Second, for travelers with high confidence in social media it is equally important to ensure that their perceived trustworthiness of the recommendation is calibrated with the actual trustworthiness. In other words, the RS should present a balanced picture of the experiences reported by guests.

CONCLUSION AND FUTURE WORK

Initial findings from the exploratory study suggest that the motivation behind searching for a recommendation influences users' decision processes. A promising idea is to investigate potential links between individual factors and presentation preferences. In contrast, the travel scenario appears to play a lesser role in personalizing the way in which recommendations are presented to users.

The user models maintained by current RS are already being exploited for personalizing the recommendation process. In addition to storing the values of various attributes (e.g., "soft bed" – important, "breakfast" – not important, "Wi-Fi" – don't care) and learned latent factors, the user model could be expanded to represent a simulation of the user's decision model. The additional user dimensions proposed in this paper could facilitate the personalization of the output. As future work, we plan to validate the proposed user model dimensions empirically using a prototype implementation built on top of an existing platform for hotel recommendations.

Personalizing the presentation of recommended items may lead to increased transparency and control over the recommendation process. Because both aspects are central to the issue of trust, we also plan to investigate whether this additional form of personalization influences the perceived trustworthiness of the 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. Betzalel, N. D., Shapira, B., & Rokach, L. "Please, not now!" A model for timing recommendations. In *Proc. RecSys '15*, ACM (2015), 297–300.
2. Bollen, D., Knijnenburg, B. P., Willemsen, M. C., & Graus, M. Understanding choice overload in recommender systems. In *Proc. RecSys '10*, ACM (2010), 63–70.
3. Bostandjiev, S., O'Donovan, J., & Höllerer, T. Tasteweights: a visual interactive hybrid recommender system. In *Proc. RecSys '12*, ACM (2012), 35–42.
4. De Carolis, B., Mazzotta, I., Novielli, N., & Silvestri, V. Using common sense in providing personalized recommendations in the tourism domain. In *Workshop on Context-Aware Recommender Systems (CARS)*, 2009.
5. Epstein, S., Pacini, R., Denes-Raj, V., & Heier, H. Individual differences in intuitive–experiential and analytical–rational thinking styles. *Journal of personality and social psychology* 71-2, (1996).
6. Jawaheer, G., Weller, P., & Kostkova, P. Modeling user preferences in recommender systems: A classification framework for explicit and implicit user feedback. *ACM Trans. Interact. Intell. Syst.* 4-2 (2014), 8:1–8:26.
7. Knijnenburg, B. P., Willemsen, M. C., Gantner, Z., Soncu, H., & Newell, C. Explaining the user experience of recommender systems. *User Mod. User-Adap.* 22, 4-5 (2012), 441–504.
8. Konstan, J. A., & Riedl, J. Recommender systems: From algorithms to user experience. *User Mod. User-Adap.* 22, 1-2 (2012), 101–123.
9. Nanou, T., Lekakos, G., & Fouskas, K. The effects of recommendations' presentation on persuasion and satisfaction in a movie recommender system. *Multimedia Systems* 16, 4-5 (2010), 219–230.
10. Pu, P., & Chen, L. Trust-inspiring explanation interfaces for recommender systems. *Knowledge-Based Systems*, 20-6 (2007), 542–556.
11. Pu, P., Chen, L., & Hu, R. Evaluating recommender systems from the users' perspective: Survey of the state of the art. In *User Mod. User-Adap.* 22, 4-5 (2012), 317–355.
12. Schafer, J. B., Konstan, J., & Riedl, J. Recommender systems in e-commerce. In *Proc. EC '99*, ACM (1999), 158–166.
13. Tintarev, N., & Masthoff, J. Designing and evaluating explanations for recommender systems. *Recommender Systems Handbook*. Springer US (2011), 479–510.
14. Verbert, K., Parra, D., Brusilovsky, P., & Duval, E. Visualizing recommendations to support exploration, transparency and controllability. In *Proc. IUI '13*, ACM (2013), 351–362.
15. Vig, J., Sen, S., & Riedl, J. Tagsplanations: explaining recommendations using tags. In *Proc. IUI '09*, ACM (2009), 47–56.