Interactive Recommending: Framework, State of Research and Future Challenges

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

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

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

EICS'16, June 21-24, 2016, Bruxelles, Belgium.

Abstract

In this paper, we present a framework describing the various aspects of recommender systems that can serve for empowering users by giving them more interactive control and transparency in the recommendation process. While conventional recommenders mostly operate like black boxes that cannot be influenced by the user, we identify four aspects properly connected with the recommendation algorithm-namely input data, user model, external context model and presentation-as essential points in which a system may be enhanced by additional interaction possibilities. In light of this framework, we take a closer look at prior and present solutions to integrate recommender systems with more inter-activity and describe future research challenges. Regarding these challenges, we especially focus on experiences gained in our own work and outline future research we have planned in the area of interactive recommending.

Author Keywords

Recommender Systems; Interactive Recommending; Models, User Experience; User Interfaces; Survey.

ACM Classification Keywords

H.3.3 [Information Storage and Retrieval: Information Search and Retrieval]: information filtering, search process; H.5.2 [Information Interfaces and Presentation: User Interfaces]: evaluation/methodology, graphical user interfaces (GUI), user-centered.

Introduction

Providing users with interactive control over the recommendation process has only recently started to receive more attention in Recommender Systems (RS) research [25, 26, 40]. In terms of objective error metrics, recommender algorithms are already quite mature and only small improvements can be expected from further optimizing algorithmic precision [40]. However, high accuracy is not the only factor determining user satisfaction [26]. It is increasingly recognized that user-related aspects such as control, trust and transparency influence the users' perception of the recommendations even more, and may contribute considerably to higher satisfaction [26, 40]. This makes it an important research goal to let users influence the recommendation process and to make it more comprehensible [25, 26, 40].

Several models exist that describe typical user behavior during the recommendation process. In earlier work [31], for instance, we have proposed a model comprising three interaction loops, which represent a) the user's interaction with the recommendations themselves, b) selection and weighting of properties related to the recommended items, and c) adaptation of entire recommender applications. Various models have also been introduced in the area of information retrieval, particularly aiming at examining the users' information-seeking behavior [28, 34]. Due to their focus on document collections and explicit search tasks, these models are however not directly applicable to RS. On the other hand, models in the area of RS research often focus on conversational and critique-based systems [44, 8], more basic feed-back processes [41], or describe system usage distinguished by different feedback types [22], i.e. ways to elicit implicit or explicit rating data. In [9], the area of interactive RS is surveyed by means of a basic model comprising those recommender components that can be extended to allow for additional interaction. While similar in some aspects to the framework we propose in this paper, the focus of the authors lies on visualizations and related aspects. How to offer users more control at the different stages in the recommendation process is only one of many aspects mentioned.

In this paper, we will therefore provide a closer look at this particular issue: First, we present a framework of interaction in RS that describes the range of possibilities users have for influencing the recommendation process. Next, we provide a detailed overview of the four aspects we have identified around the recommendation algorithm itself that allow for integrating additional interaction-input data, user model, external context model and presentation. We survey some of the most influential work related to each aspect, derive future research challenges, and outline solutions to deal with them that are especially promising from our point of view and subject of our upcoming work. Finally, we conclude the paper with a short summary and discussion.

A framework for interactive recommending

Figure 1 shows our proposed framework: Blue boxes represent components containing data, models, or presentation that may be manipulated by the user to adapt the system's outcome according to his or her current needs. The central recommender algorithm(s) (red circle) that process input data and models may also be interactively influenced, for ex-ample, by changing an algorithm's parameters or by rear-ranging the processing steps in the case of hybrid systems.

All of these components can be considered important with regard to user-perceived quality [25, 26, 40], e.g. perceived recommendation quality or transparency of the results. There have indeed been efforts to allow users to manipulate the recommender algorithms themselves [20], to

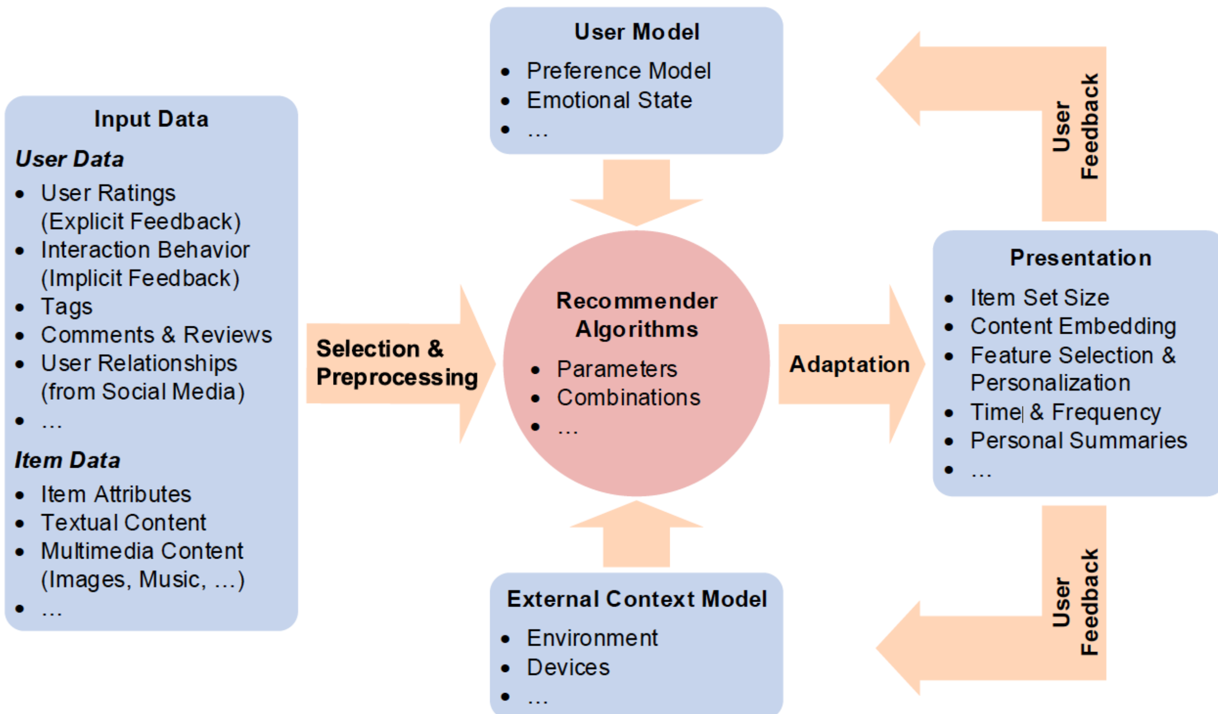

Figure 1: Framework for interactive recommending delineating the points in the recommendation process where users can be provided with additional means for interaction.

choose from different algorithms [14], or to change their influence in hybrid settings [6, 30]. However, in the following, we concentrate on a) *input data* related to users or items provided for the recommender, b) the *user model* inferred from, e.g., the user's preferences, needs, and emotions, c) the *external context model* representing the user's current situation, i.e. his or her environment, used device, etc., as well as d) the *presentation* of the recommender's results. For each aspect, a (nonexhaustive) list of properties is presented which may characterize the respective part of the system. Arrows (orange) visualize the process flow starting from possible preprocessing steps and selection of appropriate input data for the algorithms, which then generate the recommendations, i.e. adapt the presented result set. Therefore, the algorithms are able to exploit user model and external context model, which in turn may be inferred by means of the users' feedback or are generally affected by their interaction with the system.

Current position and future work

Although much effort has been put into improving the algorithms used in RS, other aspects still lack attention from the research community, especially regarding their role in increasing the recommenders' transparency and the users' influence on the systems [25, 26, 40]. In the following, we therefore have a closer look at the four relevant aspects from our model, related work and future challenges.

Input Data

The input for a RS, i.e. user or item data, is not only used by machine learning techniques to generate recommendations, but also represents an important part of such a system that might be exploited to let users influence the recommendation process and to improve their understanding of why certain items are recommended.

Collaborative Filtering (CF), the most frequently used recommender technique [42], relies on input data usually limited to user feedback, which is either explicitly provided through ratings or implicitly observed based on behavioral data [22]. Other methods use tags [43] or rely on a social graph, i.e. relationships between users [17, 18]. Particularly in content-based filtering [11], item attributes or other content- related information are used to recommend items. However, in all cases, user or item data primarily serve as input for the algorithms that generate recommendations. Only few methods exploit, for instance, tags [12, 46] or item attributes [30] to let users select and weight certain product characteristics, or visualize social connections [17] to improve users' understanding of the recommendation process.

Eliciting user preferences is an important step in order to obtain the input data necessary for the employed algorithms, which is especially relevant in cold-start situations. Various methods have been proposed to overcome the problems of traditional rating-based interfaces. Prior research has shown that ratings may be inaccurate [2] and that users prefer com-paring items instead of rating them [23]. In general, different users seem to benefit from different interaction possibilities [24]. Thus, we among others have proposed alternative preference elicitation methods: Our choice-based approach [32] allows users to state their initial preferences without the need to rate items. When compared to a conventional rating pro-cess, it has been shown to be more beneficial in terms of, e.g., perceived effort, control, and subjective recommendation quality [32]. Other authors have also experimented with novel ways to elicit preferences, for example, by letting users pick from groups of items [7] or by mapping their choice of certain pictures to factors describing their preferences [37].

We argue that exploiting input data for purposes other than feeding them into the algorithms can be an important means for giving users more control over the recommendation pro-cess. A possible challenge for future research can therefore be seen in developing techniques that create new ways of interacting with user or item data. This may comprise filtering these data even before applying the algorithms or visualizing them in order to improve the user's understanding of product space and his or her position inside it (as it has been done, for instance, through maps showing a "recommendation landscape" [16]). By building on the aforementioned works, we particularly want to improve preference elicitation for CF: Providing alternatives to simply rating a set of items seems to be a promising way to alleviate the cold-start problem [32, 37, 7, 13]. Now imagine an extension of [32] that provides users with comparisons that not directly feature the items (presented in form of, e.g., movie posters, hotel descriptions or metadata of cameras), but enables them to get an experiential impression of the products. Specifically, a system could instead use compositions of pivotal scenes captured from the movies, photos of the hotels and their amenities, or images actually taken with the respective cameras. Thus, users would be able to express their taste towards more general characteristics than just towards individual products (they may find hard to assess or do not know about).

User Model

The quality of the user model, typically learned by means of the user's feedback provided during interaction with the system, is a critical determinant for the accuracy of today's recommender algorithms. Model-based CF [42] techniques such as Matrix Factorization (MF) [27] are very prominent examples that use ratings provided by users to efficiently generate precise recommendations. The respective methods have been improved both by algorithmic advances as well as by considering additional and multiple data sources [27]. However, we argue that an adequate user model should not serve only as input for the algorithms, but might also be exploited to let users adapt the system's output and to increase their understanding of the recommendation process.

Indeed, user preferences can be modeled based on other in-puts than item ratings. In principle, all forms of implicit or explicit feedback [22], also given for item-tags [43], content-related properties, etc., can be considered. In content-based filtering, user models are typically learned by probabilistic methods or nearest neighbor algorithms based on what products the user has bought, liked or viewed before [11]. Even psychological aspects such as emotions or personality can be taken into account [39]. However, none of these approaches has been developed with the specific goal of improving interactivity. In contrast, the only way to influence the results and to (implicitly) refine the user model is typically by giving some kind of relevance feedback [11]. In social RS, it has been shown that enabling the user to adjust the importance of the mentors used for rating prediction increases transparency and satisfaction [17]. But, this is one of the only very few examples that already provide some insights in the model by means of visualizations and at the same time exploit it to allow the user actively influencing the process.

Existing interactive RS, e.g. [6, 8, 46], are often developed independently of model-based CF, and thus cannot benefit from the availability of models inferred by these efficient and accurate techniques. MF algorithms result in latent factor models where each user is individually represented by a vector whose entries describe how much the user is interested in the respective factors [27]. While it cannot be expected that improving the algorithms will further increase the actual user satisfaction with the systems [26, 40], latent factor models may also be used for other purposes than generating precise recommendations. For instance, they already have served to visualize an item landscape by reducing the high-dimensional factor space to a two-dimensional map [16]. Beyond that, the information used to model the current user's individual interests, i.e. his or her own user vector, may be exploited in even more different ways. In [38], for example, the characteristics of an item have been visualized by means of latent factors. Applying the proposed method to users instead could result in so-called 2D feature maps showing named regions that the current user is interested in. However, the only chance for users to affect their preference profile in model-based CF is usually through explicit feedback given by further ratings. In light of this fact, it is therefore-from our point of view-a major challenge to improve these systems significantly by letting users actively adjust the user model.

First attempts allow users to manipulate their user vector by other means than just rating items, i.e. more directly. With the choice-based approach mentioned before [32], it is possible to navigate through the factor space to generate a model representing the user's situational interests. Extending the landscape approach of [16] to 3D, the map's altitude can be used to reflect the user's preferences (mountains represent areas of interest while valleys indicate low relevance) [29]. In addition, the user is able to reshape the landscape in order to manipulate the user vector, thus leading to new results. We have also investigated other ways to import semantics into the abstract latent factor space, particularly by associating user-provided information such as tags with the factors [12]. While this was already known to be effective in terms of objective accuracy [27], we have confirmed this finding also with respect to subjective quality [13]. Moreover, our approach introduces a novel way to manipulate the latent user model by means of easy-to-understand tags. This seems especially useful in cold-start situations, because selecting a small number of tags leads to a meaningful new user profile without requiring the user to rate items first. Besides, as the abstract models are mostly opaque, hindering the user to understand the learned profile and hence the generated recommendations, one can imagine using the introduced semantics to better explain the user model.

Overall, while the aforementioned approaches already introduce more control over the user model, many more aspects make this part of a RS particularly interesting for increasing the level of interactivity. For example, privacy concerns suggest that users should be able to select themselves the information that will be stored in the user model and subsequently exploited for generating recommendations. Since mediating user models, i.e. importing and integrating them from other systems [4], seems promising for increasing accuracy and providing cross-domain recommendations, this should also be considered as an important subject when trying to bring more interactivity and transparency into a RS.

External Context Model

Regarding long-term interests, RS are already able to sufficiently derive the user's preferences, learn an adequate user model, and present him or her with well-fitting recommendations [26, 40, 42]. However, the user's context, i.e. date and time, season, weather, location, company of other people, used device, and many other aspects that depend on the user's current situation are often not considered in the recommendation process, although a number of context-aware recommending approaches has been proposed in recent years [1]. In fact, many systems do not even distinguish between long-term and short-term preferences, and especially disregard that the latter are strongly coupled with context [15].

A typical example is that a user might be interested in different things depending on, e.g., the currently used device: When using a smartphone on the go, he or she potentially wants suggestions for open restaurants nearby, while information that is more general would be appropriate when sit-ting in front of a desktop PC. Such variables indicated by the user's external context have already been taken into account, resulting in, among others, restaurant and travel recommenders, music recommenders specialized for different purposes (in the car, at the gym, for groups, etc.), or news RS [1]. The advent of smartphones has increased the research community's interest in developing "mobile" context-aware recommenders even more. However, although it would be particularly useful due to their increased complexity and since more information, i.e. context, has to be considered, context-aware RS often lack richer interaction possibilities [1].

So far, most work has been done on the algorithmic side, either by specializing existing methods to also integrate context or by developing techniques specifically for that use case. More details on how to incorporate contextual information may be found in [1]. However, only little attention has been paid to increasing user control in context-aware recommenders [9]. Some conversational systems adapt their dialogues implicitly based on the user's interaction sequences [33]. Similarly, changes in the user's interests can be captured to fit the results [19]. Based on the user's feedback, not only the user model, but also contextual factors can be re-fined, e.g., to filter out those restaurants that do not suit the current situation [1]. Yet overall, existing research often tries to derive the required contextual information automatically [1]. While this indeed has its benefits, letting the user actively adjust these factors is thus typically not possible-although it would give him or her the control which kind of information, e.g. about restaurants (nearby and open vs. more general), is actually desired. In [3], contextual information is used to explain recommendations, for instance, by stating that a location is especially worth a visit at a specific time of the day. In addition, the proposed system is one of the few exceptions that allows the user to influence which contextual factors to consider in the recommendation process, although this is limited to switching them on or off. Thus, finding new ways of integrating this part of a RS with interactive control seems to be a particular fruitful area of future research.

Presentation

The presentation of recommended items has also received relatively little attention by comparison. Aspects such as what information to present, how to present it, when and how often to present it, and how much of it to present for any given recommendation are important when discussing inter-activity in RS. Prior work has explored the persuasiveness of different types of recommendation lists and combinations of text with images [36]. Other researchers studied different approaches to visualize the results [45], suggested a model for timing recommendations [10], or determined the number of results that leads to high choice satisfaction without increasing choice difficulty [5]. However, most of this work stops short of considering interactivity a major factor. Consequently, ways to increase user interaction at this stage of the recommendation process remain rather unexplored.

Our work takes into consideration the recently made argument that novel approaches in RS can also stem from under-standing how people make choices [21]. Therefore, we aim to investigate choice support strategies that are not typically related to recommendation technologies, such as “combine and compute” (i.e. derive relationships from available data to show more relevant information) and “design the domain” (i.e. adapt the interface to facilitate choice) [21]. As an ex-ample, consider tourists looking for a hotel room on a booking website. Based on the choices they make during their search-destination, number of nights, desired amenities, purpose of travel, etc.-the output could be personalized not only in terms of the recommended items, but also tailored specifically to support the user’s needs. Stating a preference for “fitness center” could lead to information such as opening hours, available machines, and pricing information being displayed more prominently, or even further content being embedded, e.g. a map with related workout options nearby.

In general, a RS should be able to select the features most important for adequately personalizing the presentation ac-cording to the user’s interests and his or her situation. There-fore, the system might also leverage the wealth of information contained in user-generated data (i.e. reviews, comments, tags, or individual ratings for hotel and room characteristics) to present more relevant details about the recommended items. To illustrate this point, consider someone who is interested in venues that offer good Wi-Fi connectivity. When browsing the results, he or she might find it useful to read reviews that specifically mention aspects such as connection speed and signal strength or that give an overall quality assessment. To facilitate comparison, this information could be presented in form of a graphical scale depicting the proportion of people who rated the internet connection positively versus those who rated it negatively. Since people usually have more than one requirement, a RS that can identify the most interesting attributes for the user could enhance recommendations with such personalized summaries, thereby increasing their trustworthiness.

The presentation of results could also be improved by using social media data: By mining users’ past bookings as well as their reviews, a complex network consisting of users, hotels, and hotel attributes can be created. This would allow identifying with greater accuracy items a user is likely to find at-tractive based on the attributes mentioned in his or her re-views as well as in reviews of similar users [35]. In addition, the system could also extract and present, for each recommended item, the experiences of other people who are interested in the same attributes as the current user. Such a net-work of “co-staying in hotels” could thus introduce a novel way of increasing the interaction with RS.

Overall, as the issues mentioned before suggest, recommendations often lack transparency, and are therefore considered less trustworthy or not meeting the user’s situational needs [26, 40]. Thus, we argue that also their presentation should be adapted to better suit the current user, for example by presenting customized summaries of the recommended items as well as by identifying and selecting those features for personalization that are most important to him or her.

Conclusion

In this paper, we have summarized our experiences in the re-search area of interactive recommending. To structure the different concerns and design options for interactive RS, we presented a framework that allowed us to review the literature with respect to those aspects that bear potential for integrating the systems with additional means for interaction and may contribute to increase their transparency. For each aspect, we discussed influential existing developments in order to derive challenges for advancing the field of interactive recommending towards further improving user experience. In line with that, we also provided an outlook on some directly related future work we have planned.

References

[1] Gediminas Adomavicius and Alexander Tuzhilin. 2015. Context-aware recommender systems. In *Recommender systems handbook*. Springer, 191–226.

[2] Xavier Amatriain, Josep M Pujol, Nava Tintarev, and Nuria Oliver. 2009. Rate it again: increasing recommendation accuracy by user re-rating. In *Proceedings of the third ACM conference on Recommender systems*. ACM, 173–180.

[3] Linas Baltrunas, Bernd Ludwig, Stefan Peer, and Francesco Ricci. 2011. Context-aware places of interest recommendations and explanations. In *DEMRA '11*, Vol. 740. 19–26.

[4] Shlomo Berkovsky, Tsvi Kuflik, and Francesco Ricci. 2008. Mediation of user models for enhanced personalization in recommender systems. *User Modeling and User-Adapted Interaction* 18, 3 (2008), 245–286.

[5] Dirk Bollen, Bart P Knijnenburg, Martijn C Willemsen, and Mark Graus. 2010. Understanding choice overload in recommender systems. In *Proceedings RecSys '10*. ACM, 63–70.

[6] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. Taste-Weights: a visual interactive hybrid recommender system. In *Proceedings of RecSys '12*. ACM, 35–42.

[7] Shuo Chang, F Maxwell Harper, and Loren Terveen. 2015. Using groups of items for preference elicitation in recommender systems. In *Proceedings CSCW '15*. ACM, 1258–1269.

[8] Li Chen and Pearl Pu. 2012a. Critiquing-based recommenders: survey and emerging trends. *User Modeling and User-Adapted Interaction* 22, 1-2 (2012), 125–150.

[9] Li Chen and Pearl Pu. 2012b. Critiquing-based recommenders: survey and emerging trends. *User Modeling and User-Adapted Interaction* 22, 1-2 (2012), 125–150.

[10] Nofar Dali Betzalel, Bracha Shapira, and Lior Rokach. 2015. Please, Not Now!: A Model for Timing Recommendations. In *Proceedings of RecSys '15*. ACM, 297–300.

[11] Marco de Gemmis, Pasquale Lops, Cataldo Musto, Fedelucio Narducci, and Giovanni Semeraro. 2015. Semantics-aware content-based recommender systems. In *Recommender Systems Handbook*. Springer, 119–159.

[12] Tim Donkers, Benedikt Loepp, and Jürgen Ziegler. 2015. Merging Latent Factors and Tags to Increase Interactive Control of Recommendations. In *Poster Proceedings of RecSys âĂŹ15*.

[13] Tim Donkers, Benedikt Loepp, and Jürgen Ziegler. 2016. Tag-Enhanced Collaborative Filtering for Increasing Transparency and Interactive Control. In *Proceedings UMAP '16*. ACM, 169–173.

[14] Michael D Ekstrand, Daniel Kluver, F Maxwell Harper, and Joseph A Konstan. 2015. Letting users choose recommender algorithms: An experimental study. In *Proceedings RecSys '15*. ACM, 11–18.

[15] Mehdi Elahi, Mouzhi Ge, Francesco Ricci, Shlomo Berkovsky, and Massimo David. 2015. Interaction Design in a Mobile Food Recommender System. In *Proceedings of IntRS '15*. 49–52.

[16] Emden Gansner, Yifan Hu, Stephen Kobourov, and Chris Volinsky. 2009. Putting recommendations on the map: visualizing clusters and relations. In *Proceedings RecSys '09*. ACM, 345–348.

[17] Brynjar Gretarsson, John O'Donovan, Svetlin Bostandjiev, Christopher Hall, and Tobias Höllerer. 2010. Smallworlds: visualizing social recommendations. In *Computer Graphics Forum*, Vol. 29. Wiley Online Library, 833–842.

[18] Ido Guy. 2015. Social recommender systems. In *Recommender Systems Handbook*. Springer, 511–543.

[19] Negar Hariri, Bamshad Mobasher, and Robin Burke. 2014. Context adaptation in interactive recommender systems. In *Proceedings RecSys '14*. ACM, 41–48.

[20] F Maxwell Harper, Funing Xu, Harmanpreet Kaur, Kyle Condiff, Shuo Chang, and Loren Terveen. 2015. Putting users in control of their recommendations. In *Proceedings RecSys '15*. ACM, 3–10.

[21] Anthony Jameson, Martijn C Willemsen, Alexander Felfernig, Marco de Gemmis, Pasquale Lops, Giovanni Semeraro, and Li Chen. 2015. Human decision making and recommender systems. In *Recommender Systems Handbook*. Springer, 611–648.

[22] Gawesh Jawaheer, Peter Weller, and Patty Kostkova. 2014. Modeling user preferences in recommender systems: A classification framework for explicit and implicit user feedback. *ACM Transactions on Interactive Intelligent Systems* 4, 2 (2014), 8.

[23] Nicolas Jones, Armelle Brun, and Anne Boyer. 2011. Comparisons instead of ratings: Towards more stable preferences. In *Proceedings of WIIAT '11*, Vol. 1. IEEE, 451–456.

[24] Bart P Knijnenburg, Niels JM Reijmer, and Martijn C Willemsen. 2011. Each to his own: how different users call for different interaction methods in recommender systems. In *Proceedings RecSys '11*. ACM, 141–148.

[25] Bart P Knijnenburg, Martijn C Willemsen, Zeno Gantner, Hakan Soncu, and Chris Newell. 2012. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction* 22, 4-5 (2012), 441–504.

[26] Joseph A Konstan and John Riedl. 2012. Recommender systems: from algorithms to user experience. *User Modeling and User-Adapted Interaction* 22, 1-2 (2012), 101–123.

[27] Yehuda Koren, Robert Bell, Chris Volinsky, and others. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.

[28] Carol C Kuhlthau. 1991. Inside the search process: Information seeking from the user's perspective. *J. Am. Soc. Inf. Sci.* 42, 5 (1991), 361.

[29] Johannes Kunkel, Benedikt Loepp, and Jürgen Ziegler. 2015. 3D-Visualisierung zur Eingabe von Präferenzen in Empfehlungssystemen [3D visualization to elicit preferences in recommender systems]. *Proceedings of M&C âĂŽ15* (2015).

[30] Benedikt Loepp, Katja Herrmanny, and Jürgen Ziegler. 2015a. Blended recommending: Integrating interactive information filtering and algorithmic recommender techniques. In *Proceedings of CHI '15*. ACM, 975–984.

[31] Benedikt Loepp, Katja Herrmanny, and Jürgen Ziegler. 2015b. Merging Interactive Information Filtering and Recommender Algorithms–Model and Concept Demonstrator. In *icom*, Vol. 14. 5–17.

[32] Benedikt Loepp, Tim Hussein, and Jüergen Ziegler. 2014. Choice-based preference elicitation for collaborative filtering recommender systems. In *Proceedings of CHI '14*. ACM, 3085–3094.

[33] Tariq Mahmood, Francesco Ricci, and Adriano Venturini. 2009. Improving recommendation effectiveness: Adapting a dialogue strategy in online travel planning. *J. of IT & Tourism* 11, 4 (2009), 285–302.

[34] Gary Marchionini. 1997. *Information seeking in electronic environments*. Number 9. Cambridge University Press.

[35] Julian McAuley and Jure Leskovec. 2013. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings RecSys '13*. ACM, 165–172.

[36] Theodora Nanou, George Lekakos, and Konstantinos Fouskas. 2010. The effects of recommendationsâĂŹ presentation on persuasion and satisfaction in a movie recommender system. *Multimedia systems* 16, 4-5 (2010), 219–230.

[37] Julia Neidhardt, Rainer Schuster, Leonhard Seyfang, and Hannes Werthner. 2014. Eliciting the users' unknown preferences. In *Proceedings RecSys '14*. ACM, 309–312.

[38] Bottyán Németh, Gábor Takács, István Pilászy, and Domonkos Tikk. 2013. Visualization of movie features in collaborative filtering. In *Proceedings of SoMeT '13*. IEEE, 229–233.

[39] Maria Augusta SN Nunes and Rong Hu. 2012. Personality-based recommender systems: an overview. In *Proceedings RecSys '12*. ACM, 5–6.

[40] Pearl Pu, Li Chen, and Rong Hu. 2012. Evaluating recommender systems from the userâĂŹs perspective: survey of the state of the art. *User Modeling and User-Adapted Interaction* 22, 4-5 (2012), 317–355.

[41] Pearl Pu, Boi Faltings, Li Chen, Jiyong Zhang, and Paolo Viappiani. 2011. Usability guidelines for product recommenders based on example critiquing research. In *Recommender Systems Handbook*. Springer, 511–545.

[42] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2015. Recommender systems: introduction and challenges. In *Recommender Systems Handbook*. Springer, 1–34.

[43] Shilad Sen, Jesse Vig, and John Riedl. 2009. Tagommenders: connecting users to items through tags. In *Proceedings of WWW '09*. ACM, 671–680.

[44] Barry Smyth and Lorraine McGinty. 2003. An analysis of feedback strategies in conversational recommenders. In *Proceedings of AICS '03*.

[45] Katrien Verbert, Denis Parra, Peter Brusilovsky, and Erik Duval. 2013. Visualizing recommendations to support exploration, transparency and controllability. In *Proceedings of IUI '13*. ACM, 351–362.

[46] Jesse Vig, Shilad Sen, and John Riedl. 2011. Navigating the tag genome. In *Proceedings of IUI '11*. ACM, 93–102.