User Control in Recommender Systems: Overview and Interaction Challenges

Dietmar Jannach, Sidra Naveed, and Michael Jugovac

Department of Computer Science
TU Dortmund, Germany
{firstname.lastname}@tu-dortmund.de

Abstract. Recommender systems have shown to be valuable tools that help users find items of interest in situations of information overload. These systems usually predict the relevance of each item for the individual user based on their past preferences and their observed behavior. If the system's assumption about the users' preferences are however incorrect or outdated, mechanisms should be provided that put the user into control of the recommendations, e.g., by letting them specify their preferences explicitly or by allowing them to give feedback on the recommendations. In this paper we review and classify the different approaches from the research literature of putting the users into active control of what is recommended. We highlight the challenges related to the design of the corresponding user interaction mechanisms and finally present the results of a survey-based study in which we gathered user feedback on the implemented user control features on Amazon.com.

1 Introduction

Recommender systems have become an integral part of many commercial websites like Amazon, Netflix, and YouTube. In scenarios where millions of choices are available these systems serve as an aid for users in their search and decision making processes by automatically assessing the users' preferences and by making personalized recommendations.

On many websites, including the above-mentioned ones, the underlying user preference model is established by the system by observing and interpreting the users' behavior over time ("implicit feedback") or by considering the user's explicit ratings for individual items. The estimated preference model is then used to make predictions about the relevance of each recommendable item for the user. Over the last two decades a variety of algorithms was proposed in the literature to optimize these relevance assessments using datasets that represent a snapshot of the user's preferences.

In reality, however, the user interests can change over time, which means that some preference information can become outdated, leading to inaccurate recommendations [1, 2]. Furthermore, the relevance of an item can depend on the user's current situation. The user might, for example, be looking for a gift which does not fall into his or her typical preference profile. Or, the user might have just bought a certain item so that further recommending the same or similar objects becomes pointless.

In many real-world recommender systems users have limited or no means to inform the system that its assumptions are incorrect or to specify that preference information has become outdated.[1] Past research has however shown that at least in some application domains users appreciate being more actively involved in the process and in control of their recommendations [3, 4]. In the end, providing additional forms of user interactions can not only lead to higher user satisfaction but also increase the users' trust in the system [5, 6].

In the research literature, a number of proposals have been made on how to implement mechanisms for increased user control. Simple approaches are, for example, based on static preference forms. Others use conversational dialogs or critiquing mechanisms to let the users specify their constraints and preferences. Some proposals even allow the user to choose between different recommendation strategies. Generally, the proposed mechanisms provide different levels of user control but they unfortunately all come with their own challenges regarding the user interaction.

In this paper we first provide an overview of user control mechanisms from the literature, categorize them according to their context in the recommendation process, and discuss the individual user interaction challenges. As a case study of a real system, we then report the findings of a survey-based study in which we investigated how users perceive the comparably powerful explanation, feedback, and control mechanisms that are implemented on Amazon's website. Our observations indicate that although the implemented features are known to many study participants, most users are hesitant to use the provided functionality for different reasons.

2 Control Mechanisms and User Interaction Challenges

2.1 Conceptual Framework

Overview. Figure 1 shows an overview of the approaches and situations in the recommendation process where users can be put into control according to the literature. We categorize the different techniques in two classes:

- Techniques where users are allowed to explicitly specify their preferences. These will be discussed in Section 2.2.
- Techniques that put the user into control in the context of recommendation results. We review these approaches in Section 2.3.

Critiquing-based techniques share characteristics of both categories. We will discuss them also in Section 2.2.

[1] In some rating-based systems users can update their ratings, which might however be tedious, and changes often have no immediate effect on the presented recommendations.

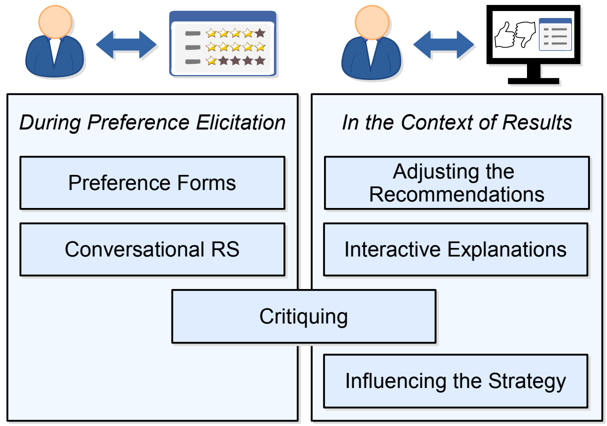

Fig. 1: User Control Mechanisms for Recommender Systems

Definition of User Control. In the context of this work we require that user control mechanisms have an *immediate effect* on the recommendations. For example, selecting a certain interest category in a preference form should immediately lead to updated results the next time the recommendations are displayed. Changing or adding explicit item ratings therefore do not count as control mechanisms as usually these changes are not immediately reflected in the results, e.g., because the trained models are only updated periodically. The same holds for like/dislike buttons, which some websites display for each recommendation, in case these have no immediate effect on the next recommendations.

Our second, but softer, requirement is that users should understand or at least have a confident intuition about the effects of their control actions. A "thumbs down" action for the currently played track on the music service Spotify for example results in an immediate update of the next tracks to be played. The logic behind the update is however not transparent, which is why we consider this as a limited form of user control.

Finally, control in recommender systems is sometimes discussed in the literature together with "inspectability", e.g., in [3]. Inspectability, i.e., giving the user insights on what the recommendations are based on, is in our view not a requirement for a control mechanism but can be useful to help users understand the possible effects of their control actions.

2.2 Control During the Preference Elicitation Phase

Preference Forms and Adaptive Dialogs One basic option of giving control to the users is to let them specify their constraints and preferences explicitly by using *static user profile forms*. Figure 2a shows a simple screen that allows users to choose their genre interest for the Netflix movie recommender. In some applications, such preference forms are used to indirectly infer the interests, e.g., by asking the users for their favorite movies or artists. Such simple forms of user control during preference elicitation are for example implemented in the music

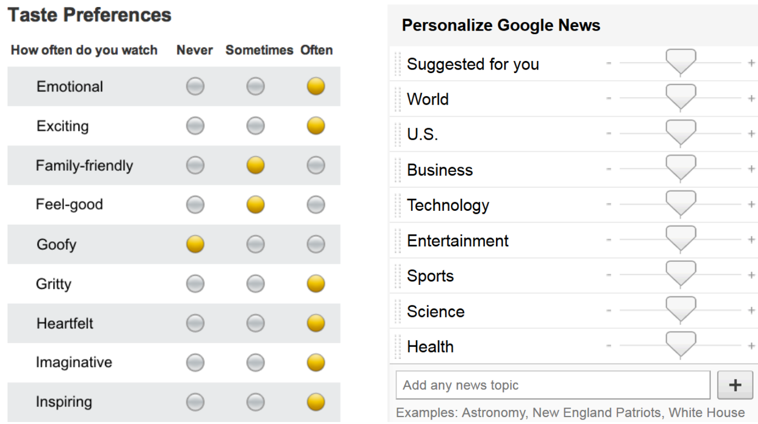

(a) Taste Preference Selection on the Netflix Movie Streaming Site

(b) Preference Indicators on the Google News Site

Fig. 2: Static Preference Forms for Personalized Services

recommender presented in [7], in the MenuMentor restaurant menu system [8], and in the energy-saving application described in [9]. A similar approach is also implemented on Google News, where users indicate their preferences about news in different categories using slider controls (see Figure 2b).

Generally, static forms are comparably easy to use. However, as soon as the user is allowed to indicate *relative preferences*, these forms can become complicated in terms of their interpretation. For example, in case of the Google News preference indicators it is not clear if having all sliders in the middle position has the same meaning as having all at the maximum level. Another problem with such static forms is that every time the users' interests change, they have to manually adapt their settings such that they properly reflect their new interests.

Because static forms are identical for all users, they might not be optimally suited to capture the preferences of all kinds of users, who can have different levels of expertise in a domain. *Conversational approaches* in some sense try to mimic a human advisor for high-involvement products like digital cameras, e.g., by guiding the user through an interactive dialog based on desired functional features or by providing additional explanations when requested. An early system is the ADAPTIVE PLACE ADVISOR [10], which, according to the classification in [11], adapts its conversation behavior to the users at the information filtering and navigation levels. Similar ideas are implemented in the ADVISOR SUITE system [12], which also adapts the conversation based on the user's previous answers and in addition is capable of explaining the recommendations and can help users in situations in which no item fulfills all their requirements.

Technically, these conversational systems often implement item filtering rules that deterministically map functional features to technical product characteristics. Users of such systems are therefore in immediate control of the recommendation outcomes. Implementing such systems can however require significant knowledge engineering efforts to encode and maintain the recommendation rules. Usually, these systems also do not learn over time or from the behavior of a larger community. From an interaction perspective, users can also feel overwhelmed when they try to change some of their specifications after the initial conversational elicitation phase.

Critiquing Similar to the discussed form-based techniques, users of *critiquing* approaches explicitly state their preferences on certain item features. Here, however, they do that in the context of a reference item, e.g., a camera, and the provided preferences are *relative* statements like “cheaper” or “higher resolution”. The system then uses this feedback to find other items that fulfill the refined requirements. The process is repeated until the user finds a suitable camera or gives up on the search. Critiquing based systems were presented, e.g., in [8, 13, 14], and a number of works have been proposed to improve the basic interaction scheme, including *compound* or *dynamic* critiques, where users can for example update their preferences in more than one dimension at a time.

Critiquing approaches have the advantage that their general operating principle is easy to understand for the users. Furthermore, each interaction is followed by an immediate update of the recommendation(s). However, basic critiquing schemes can lead to a high number of required iterations until a suitable product is found. Compound critiques, on the other hand, can induce higher cognitive load for the users. Finally, similar to form-based approaches the problem can arise that no more suitable items remain that can be recommended, which means that the system has to implement a recovery strategy for the user.

2.3 Control in the Context of Recommendation Results

Dynamically Adjusting the Recommendations Once a set of recommendations is computed, a simple form of allowing users to influence what is presented is to provide them with mechanisms to further filter and re-sort the items based on their features. Such a post-filtering functionality was for instance implemented in the MovieCritic [15] and the MetaLens systems [16], where users could for example include or exclude movies of certain genres. In the MetaLens system, users could also indicate the relative importance of individual features.

A more sophisticated and visually complex approach was proposed in [17]. Their system displays three pieces of information in parallel – the items that the user has rated, the corresponding set of similar users, and the resulting recommendations. Users can then not only inspect why certain items were recommended but also interactively adapt their ratings, which is then reflected in updated recommendations.

TasteWeights [18] is a similar approach that also combines a visualization of the recommendation logic with an interaction mechanism. Their system presents a graph that shows the relationships between the user’s rated items, their social friends, and the recommendations. The implemented control mechanism allows users to adjust the weights of the items and the friends. Similar to the work in [17], an experimental evaluation indicates that such a form of user involvement can lead to higher user satisfaction.

Another comparable approach was proposed in [19], where a web-based interactive visualization for a content recommender system for microblogs was devised. Their interface also consists of three columns and users can for example change the sort criterion of the items (tweets) or vary the relative importance of different filters.

Overall, all of the presented approaches to put users into control lead to immediate effects on the resulting recommendations. In most cases, the users will at least to some extent understand how their actions (indirectly) impact the outcomes. However, one cannot generally assume that average users will understand the underlying rationale of, e.g., a neighborhood based method. A limitation of some of the works is in fact that they employ comparably simple recommendation methods, and it is unclear how such approaches would work for more complex machine learning models. In addition, users might have rated dozens of items over time and might have a large neighborhood so that manipulating ratings or weights on such a fine-grained level might soon become tedious.

User Control in the Context of Explanations The literature suggests that providing explanations for recommendations can be beneficial in different ways as they, for example, help users understand why certain items were recommended. This in turn could lead to increased satisfaction and trust in the system [20]. Since mechanisms for user control often require that users understand or at least have an intuition of the reasoning logic of the system, designing these mechanisms in the context of explanations appears natural.

In the context of conversational systems, such interactive explanations were for example developed for the Advisor Suite system described above [21]. The knowledge-based system was able to generate textual explanations based on the logical rules that map user preferences to item characteristics. In case some requirements could not be fulfilled, users were able to overwrite the default priorities of the rules with their personal preference weights. This feedback was then immediately processed to compute an updated recommendation list.

Feedback mechanisms in the context of explanations were also implemented in the mobile recommender system Shopr [22]. In this system the explanations were provided along with the recommendations, e.g., *"Because you were interested in … in the past"*, and users could then give feedback to the system about whether this recommendation logic was inappropriate in their current situation. A possible type of feedback included not only to indicate that a certain item is not relevant but a whole category of items should not be recommended. A similar feature, although not in the context of explanations, can be found on YouTube, see Fig. 3.

The discussed methods in general allow users to correct possibly wrong assumptions in the context of what is sometimes called "scrutable" interactive explanations [23]. The concept of these scrutable explanations is that with their help users are able to inspect and understand (scrutinize) the system's reasoning and act upon this knowledge to improve the system's assumptions [20].

As with all forms of user control discussed so far, it can be challenging to design such interactive explanations when the underlying reasoning mechanisms are complex. In these cases, generating understandable explanations can represent

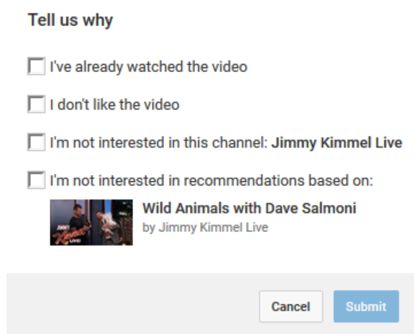

Fig. 3: Feedback Options for a Recommendation on the YouTube platform.

a problem of its own. In Section 3, we will discuss the interactive explanation mechanism that is implemented on Amazon.com in more detail.

Choosing or Influencing the Recommendation Strategy A quite different approach of letting users influence the recommendations is to allow them to select or parameterize the algorithms themselves that are used to generate the recommendations. In the study described in [24], for example, users of the MovieLens system were able to choose one of four predefined algorithms by clicking on a widget in the top menu bar. Each selection immediately led to a different set of recommendations. An analysis of the log files revealed that about one quarter of the users actually tried out the recommendation-switching feature.

More fine-grained control was given to users in the approach presented in [25], where users could fine-tune the importance weights of a hybrid recommender. Their graphical interface furthermore visualized – with the help of a Venn diagram – based on which algorithm each item was included in the recommendations. Two user studies were performed to assess the effect of the system on the users and the authors report that their system led to higher user engagement, and it furthermore seems that users worked more efficiently with the tool.

User control in terms of interactively fine-tuning the desired item characteristics was recently proposed and experimentally analyzed in [26]. The participants of a user study could for example change the popularity or recency level of the movies to be recommended. When the results were presented, the users could then use additional knobs to fine-tune the results until they were satisfied. An analysis revealed that, in the end, users were happier with the adapted recommendations than with the original ones.

Overall, the different studies of letting users control the underlying algorithms indicate that such mechanisms can have a positive effect on the user experience. Some of the proposed methods are however comparably simple. Letting users vary the popularity of the recommended movies can be seen as a form of changing the sort order. To some extent it therefore remains unclear how user interfaces should be designed for more complex fine-tuning functionalities as they have to be intuitive and easy to use for a heterogeneous user community.

Fig. 4: Explanation-based Control Mechanism at Amazon.com

3 On the Acceptance of Amazon's Scrutable Explanations

Our literature overview in general indicates that finding appropriate user interface mechanisms for putting users into control can be challenging. In the end, a poor UI design can lead to limited acceptance of the control mechanisms by the users, e.g., because they find the system tedious or have problems understanding the effects of their actions.

To obtain a better understanding of mechanisms for user control in the context of recommendations, we conducted a user survey about the acceptance of the comparably rich explanation, feedback, and control functionality that is implemented on the websites of Amazon. Specifically, our goals were to assess to what extent users are aware of the functionality, if they find the provided functionality clear, and if they are actually using it.

User feedback and control on Amazon.com is provided in the context of explanations as proposed in the literature in [21, 22], or [23]. Figure 4 shows a screenshot of the provided functionality. The presented screen can be reached by navigating to dedicated pages on which users can inspect and improve their recommendations. For some product categories the functionality can be accessed directly from the recommendation lists, but Amazon seems to vary the UI design in that respect over time.

In Figure 4 a recommended item is displayed along with an explanation why it is recommended. The explanation in this case refers to another item that the user already owns. The user can then give feedback on the recommended item by rating it or indicating no interest in the item. Furthermore, the users can instruct the system not to use the already owned item for future recommendations. This functionality can for example be helpful when the user's interest has changed over time or when the item is outside of his or her usual interests, e.g., because it was a gift.

3.1 Survey Design

We conducted a two-phase survey-based study in which we asked users of Amazon.com about their knowledge and usage behavior regarding the feedback mechanisms implemented on the site. All participants were computer science students of our university and at the same time regular customers of Amazon doing several purchases a year.

Phase 1: 75 students participated in the first part. The survey sheet showed screenshots of Amazon's functionality and comprised 15 questions with 5-point Likert-scale answer options. Questions were for example about whether or not the participants know the functionality, if they find the functionality clear, and whether or not they have used or intend to use it in the future.[2]

Phase 2: The second phase, which took place a few weeks later, particularly focused on reasons why the users would or would *not* click on the recommendations and on possible reasons for *not using* the feedback functionality. 26 students of the user population from the first phase returned the survey sheets, which – besides a set of Likert-scale questionnaire items – included two free-text fields where the participants could express reasons *not to use* the provided functionalities.

3.2 Observations

Phase 1. A first surprising observation is that more than 75% of the participants stated that they use recommendations on the site never or only rarely. At the same time, they found it on average "rather clear" or "very clear" how the recommendations are created. The average answer value was 4.04 on the five-point scale, where five means "very clear".

When asked whether they knew that they could influence their recommendations, more than 90% of the subjects answered positively, even though not all of them knew exactly how. However, only about 20% were aware of the special page for improving recommendations and even fewer had ever used the page.

Regarding the feature called *"Don't use for recommendations"*, more than 50% stated that the provided functionality was clear or very clear to them. Another 25% said that one could guess its purpose. On the other hand, only about 8% of the users (6 out of 76) had ever actually used the functionality.

We then asked the participants whether they had ever noticed the provided explanation (*"Because you bought ..."*). Around 40% answered with yes. Again, the majority of subjects stated that the functionality is mostly clear to them, but only 4 of 76 had ever used the *"Don't use for recommendation"* feature in that context. Finally, although the participants found the functionality potentially useful (avg. answer was 3.36 on the five-point scale, where five represents the most positive answer), the intention to use the feature in the future was rather limited (2.55).

Phase 2. In the second phase we were particularly interested in the reasons why the participants are hesitant to use the provided recommendation, feedback

[2] A translated version of the survey forms can be found at `http://ls13-www.cs.tu-dortmund.de/homepage/publications/ec-web-2016/`

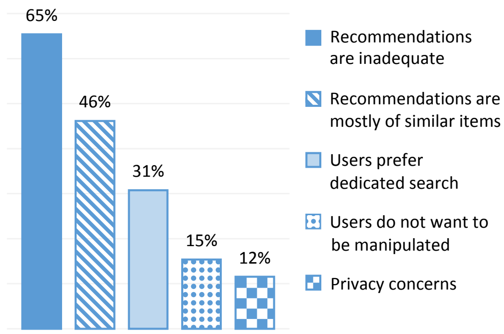

(a) Why Participants' Do Not Use Recommendations

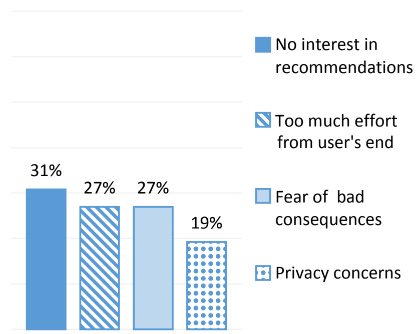

(b) Why Participants' Do Not Use Control Mechanisms

Fig. 5: Results of the Amazon Survey

and control mechanisms. We manually analyzed the qualitative free-form feedback from the subjects and categorized the responses in different groups.

Figure 5a shows the statistics for the reasons why many of the participants would not use Amazon's recommendations. The two main reasons are related to the quality of the recommendations, which appear not to be very helpful or contain too many similar items.[3] Another common reason for not relying on recommendations is that users prefer to use explicit search. Finally, privacy concerns and fear of being manipulated were other aspects mentioned by the participants in this survey.

In Figure 5b we summarize the reasons for not using the *"Don't use for recommendations"* feature. One main reason is that the participants do not use the recommendations in the first place. Many however also found that this form of fine-tuning requires too much effort. An equal amount of respondents were afraid of the consequences of their actions and of the inability to undo their settings later on. A smaller amount of participants again mentioned privacy issues.

3.3 Discussion

Our results indicate that although the mechanisms provided on Amazon.com were known to many participants, they are hesitant to actually use them, e.g., due to the additional effort or unclear consequences of their actions.

Furthermore, the responses of the users in general indicate limited satisfaction and trust in the recommendation and feedback system. Providing mechanisms that are understandable for users and have an immediate effect on the recommendations seems to be required but not sufficient, which calls for better mechanisms to put users into control. More user-friendly systems could, for example, provide less tedious forms of interaction or clearly indicate that profile changes can be undone to reduce the users' fear of undesired consequences.

[3] The participants could provide several reasons and the value 65% indicates that nearly two thirds of the users stated that the recommendations were inadequate.

Overall, we see our survey as a further step toward a better understanding of user control mechanisms for recommenders. However, a number of questions remains open for future research. Further studies could, for example, continue related lines of research described in [9, 27, 28, 29, 30] and further investigate what level of control users expect from recommender systems, whether more control is always better, or if different users call for different control mechanisms. Also, further research is necessary to identify the effects of control mechanisms on user engagement and the user experience.

Regarding research limitations, note that our survey is based on the responses of computer science students, who might have a representative online shopping behavior for their age group but are maybe untypical in some aspects, e.g., with respect to privacy concerns. The sample size of the initial study reported in this paper also represents a research limitation.

4 Conclusions

We have reviewed the literature on user control in recommender systems and have identified different requirements to make such approaches effective in particular with respect to the design of the user interaction mechanisms. A survey among users of Amazon.com indicates that the provided functionality is only used to a very limited extent. Besides the poorly regarded quality of the recommender system itself, the major reasons include the partially unclear consequences of feedback and control actions and the increased user effort, which indicates that more research is required in this area.

References

1. Hu, Y., Koren, Y., Volinsky, C.: Collaborative filtering for implicit feedback datasets. In: ICDM '08. (2008) 263–272
2. Amatriain, X., Pujol, J.M., Tintarev, N., Oliver, N.: Rate it again: Increasing recommendation accuracy by user re-rating. In: RecSys '09. (2009) 173–180
3. Knijnenburg, B.P., Bostandjiev, S., O'Donovan, J., Kobsa, A.: Inspectability and control in social recommenders. In: RecSys '12. (2012) 43–50
4. Dooms, S., De Pessemier, T., Martens, L.: Improving IMDb movie recommendations with interactive settings and filter. RecSys '14 (2014)
5. McNee, S.M., Lam, S.K., Konstan, J.A., Riedl, J.: Interfaces for eliciting new user preferences in recommender systems. In: UM '03. (2003) 178–187
6. Xiao, B., Benbasat, I.: E-Commerce product recommendation agents: Use, characteristics, and impact. MIS Quarterly **31(1)** (2007) 137–209
7. Hijikata, Y., Kai, Y., Nishida, S.: The relation between user intervention and user satisfaction for information recommendation. In: SAC '12. (2012) 2002–2007
8. Wasinger, R., Wallbank, J., Pizzato, L., Kay, J., Kummerfeld, B., Böhmer, M., Krüger, A.: Scrutable user models and personalised item recommendation in mobile lifestyle applications. In: UMAP '13. (2013) 77–88

9. Knijnenburg, B.P., Reijmer, N.J., Willemsen, M.C.: Each to his own: how different users call for different interaction methods in recommender systems. In: RecSys '11. (2011) 141–148
10. Goker, M., Thompson, C.: The Adaptive Place Advisor: A conversational recommendation system. In: 8th German Workshop on CBR. (2000) 187–198
11. Kobsa, A., Koenemann, J., Pohl, W.: Personalized hypermedia presentation techniques for improving online customer relationships. The Knowledge Engineering Review **16**(2) (2001)
12. Felfernig, A., Friedrich, G., Jannach, D., Zanker, M.: An integrated environment for the development of knowledge-based recommender applications. International Journal of Electronic Commerce **11**(2) (2006) 11–34
13. Burke, R.D., Hammond, K.J., Young, B.C.: Knowledge-based navigation of complex information spaces. In: AAAI '96. (1996) 462–468
14. Trewin, S.: Knowledge-based recommender systems. Encyclopedia of Library and Information Science **69** (2000) 180–200
15. Swearingen, K., Sinha, R.: Beyond algorithms: An HCI perspective on recommender systems. In: ACM SIGIR Recommender Systems Workshop. (2001) 1–11
16. Schafer, J.B., Konstan, J.A., Riedl, J.: Meta-recommendation systems: user-controlled integration of diverse recommendations. In: CIKM '02. (2002) 43–51
17. Schaffer, J., Höllerer, T., O'Donovan, J.: Hypothetical recommendation: A study of interactive profile manipulation behavior for recommender systems. In: FLAIRS '15. (2015) 507–512
18. Bostandjiev, S., O'Donovan, J., Höllerer, T.: Tasteweights: a visual interactive hybrid recommender system. In: RecSys '12. (2012) 35–42
19. Tintarev, N., Kang, B., Höllerer, T., O'Donovan, J.: Inspection mechanisms for community-based content discovery in microblogs. RecSys IntRS '15 Workshop (2015) 21–28
20. Tintarev, N., Masthoff, J.: A survey of explanations in recommender systems. In: IEEE ICDEW Data Engineering Workshop. (2007) 801–810
21. Jannach, D., Kreutler, G.: Rapid development of knowledge-based conversational recommender applications with Advisor Suite. Journal of Web Eng. **6(2)** (2007)
22. Lamche, B., Adıgüzel, U., Wörndl, W.: Interactive explanations in mobile shopping recommender systems. In: RecSys IntRS '14 Workshop. (2014) 14–21
23. Czarkowski, M., Kay, J.: A Scrutable Adaptive Hypertext. In: AH' 02. (2002) 384–387
24. Ekstrand, M.D., Kluver, D., Harper, F.M., Konstan, J.A.: Letting users choose recommender algorithms: An experimental study. In: RecSys '15. (2015) 11–18
25. Parra, D., Brusilovsky, P., Trattner, C.: See what you want to see: visual user-driven approach for hybrid recommendation. In: IUI '14. (2014) 235–240
26. Harper, F.M., Xu, F., Kaur, H., Condiff, K., Chang, S., Terveen, L.G.: Putting users in control of their recommendations. In: RecSys '15. (2015) 3–10
27. Jameson, A., Schwarzkopf, E.: Pros and cons of controllability: An empirical study. In: AH' 02. (2002) 193–202
28. Kramer, T.: The effect of measurement task transparency on preference construction and evaluations of personalized recommendations. Journal of Marketing Research **44(2)** (2007) 224–233
29. Chen, L., Pu, P.: Critiquing-based recommenders: survey and emerging trends. User Modeling and User-Adapted Interaction **22(1-2)** (2012) 125–150
30. Groh, G., Birnkammerer, S., Köllhofer, V.: Social recommender systems. In: Recommender Systems for the Social Web. Springer (2012) 3–42