

Chapter 5

Handling Preferences



Alexander Felfernig and Martijn Willemsen

Abstract This chapter presents an overview of approaches related to the handling of preferences in (group) recommendation scenarios. We first introduce the concept of *preferences* and then discuss how preferences can be handled for different recommendation approaches. Furthermore, we sketch how to deal with inconsistencies such as contradicting preferences of individual users.

5.1 Introduction

Before making recommendations, it is necessary to know and understand the preferences of the users you are trying to serve [16]. Recommender systems create different types of preference models in order to discern the relevance of items. The term *preference* in recommender systems can be loosely characterized as *something that refers to the things in a user's head that determine how he/she will evaluate particular alternatives, and what choices he/she will make* [38, 40]. In this broad sense, preferences refer either to taste or to the utility of items (e.g., I like strawberry ice cream), or to the outcome of a decision process: I *prefer* strawberry over chocolate ice cream. In this latter sense, *preference* is by nature a relative statement. As discussed in De Gemmis et al. [16], a preference can also be regarded as an ordering relation between two or more items to describe which of a given set of alternatives best suits a user. Jameson et al. [38] differentiate between *general* and *specific preferences* where the former is related to evaluations on a categorical level¹ (e.g., *economy* of a car is more important than *sportiness*) and the latter to items or attributes (e.g., I prefer to see the movie *Transformers IV* over *Transformers V*).

¹Level of interest dimensions [75].

Acquiring the preferences of users and interpreting these in a way that leads to items relevant for users is often a difficult task [16]. Traditional microeconomic models of human decision making assume that consumers are able to make optimal decisions [33, 52]. These models assume that human preferences are the result of a formal process of utility maximization where item utilities and attributes are fully known and remain stable over time. In many real-world settings, this assumption does not hold. For example, if a family wants to purchase a new car, an upper price limit could have been defined at the beginning of the decision process but then be revised in the face of new highly relevant features that were not considered beforehand. Preferences can change because our utilities for items or features change due to the context of the task [5], or simply because relevant features only come to mind over the course of the decision making process.

This evidence against the assumption of a given set of stable preferences led to alternative models of human decision making [56, 59, 65] and also coined the term *preference construction* [5, 44] which states that in many decision making situations, people construct their specific preferences for options while making the decision. In one way or another, most existing recommender systems take into account the fact that preferences are strongly influenced by user goals, personal experiences, information received from family and friends, and cognitive limitations [38]. Depending on the recommendation approach, specific aspects are taken more into account than others. For example, critiquing-based recommendation approaches take into account a user's limited knowledge about the item domain in terms of supporting the exploration of the search space on the basis of critiques; collaborative recommendation approaches simulate recommendations received from family and friends, but still assume that preferences for items, as reflected in their ratings, are stable, like traditional economic models do.

User feedback regarding specific preferences can be given in an explicit (the user is "actively" involved in the elicitation task) or implicit fashion (the user is not "actively" involved) [16, 59] (see Table 5.1). *Explicit feedback* is given, for example, by rating choice alternatives (*relevance feedback*) [16], critiquing the currently presented reference item [9, 64], ranking options via pairwise preferences [41] or choice-based preference elicitation [32], and in terms of explicit preferences with regard to item properties (specifically, in constraint-based recommendation scenarios) [37]. The advantage of explicit methods is an explicit link between the feedback given and the preference that is measured, but this comes at the disadvantage of requiring effort and the active involvement of users, which is not always practical in real life applications. Therefore, recommender systems often use *implicit feedback* that can be collected by observing a user's navigation and purchasing behavior. Implicit feedback is also given in terms of a user's eye movements when interacting with a recommender system [77], movements of users in public contexts [43], or a user's item purchases [15]. However, the link between the user's behavior and the specific preferences and goals of the user is only indirect. There are limits as to what can be inferred through observation [18].

In this chapter, we analyze preference elicitation support in different recommendation approaches (collaborative filtering, content-based filtering, constraint-based,

Table 5.1 Representation of user preferences (see also [58–60])

Recommendation approach	Explicitly formulated preferences	Implicitly formulated preferences
Collaborative filtering	Item ratings [19], pairwise preferences [41], choice-based preference elicitation [32]	Item reviews [10], user location data [61], time of item consumption [72]
Content-based filtering	Item ratings, categories and tags [57], excluded items [8]	Extracted keywords [57], eye movements [77], item reviews [10]
Constraint-based (incl. utility-based) recommendation	Attribute values [22], preferences between attribute values [7, 39, 71], attribute weights [21, 46], interest dimensions [22, 54]	Items selected for comparison, degree of domain knowledge derived from induced conflicts [23]
Critiquing-based recommendation	Critiques on item attributes [51], natural language based critiques [31]	Information from chats [55], eye movements [11]

and critiquing-based recommendation) [26] and also discuss specific aspects related to the group context. Furthermore, we point out ways to deal with inconsistencies in a given set of user preferences [25].

5.2 Collecting Preferences

Depending on the recommendation approach, preferences are observed / collected in different ways. An overview of the different types of preference representations used in recommendation scenarios is given in Table 5.1. Most *group recommender applications* apply preference elicitation approaches that are quite similar to approaches in single user recommender systems [2, 36]. Where appropriate, we will point out relevant differences.

Preferences in Collaborative Filtering The dominant approach to providing explicit preference feedback is to rate items [19, 76]. Implicit preferences are given in the form of item reviews, user location data, and point of time of item consumption [10, 61, 72]. In the context of collaborative-filtering-based group recommender systems, the individual assessments of items represent the (sometimes aggregated) preferences of individual group members. In this context, typically N-point response scales (e.g., 5-star rating scales) are used to represent user feedback. Different rating scales are used in collaborative filtering recommender systems, for example, the MOVIELENS recommender system [53] offers a 5-point scale (with half-star ratings) whereas the JESTER *joke recommender system* provides a continuous rating scale between -10 and $+10$. LAST.FM provides a binary rating scale and NETFLIX recently switched to a thumbs up/down rating, replacing its 5-star rating scale as A/B tests showed it increased explicit user feedback by 200%. This shows that

scale granularity reflects a tradeoff between cognitive effort [68] and amount of information acquired [42]. As ratings only provide an evaluation of solitary items [38], methods have been proposed that (1) take into account pairwise preferences that measure the relative preference between two items [41] or (2) elicit user preferences from list representations, adaptively changing the list to gradually discover the user's preference [32]. Though these alternative methods have not been applied directly to group recommendations, one can envision that asking a group to rank items rather than rate them might provide a more efficient and satisfactory way to discover a ranking that best fits the preferences of the entire group.

Preferences in Content-Based Recommendation Explicit preference feedback is provided in the form of item evaluations and the specification of meta-properties represented, for example, as categories or tags [57]. Implicit preferences are specified, for example, in terms of item reviews [10] and eye movement patterns (collected via eye-tracking) [77]. In the context of content-based group recommender systems, ratings and category preferences represent the preferences of individual group members. As pointed out, for example, in [8], it often makes sense to explicitly specify and represent negative preferences. Taking such information into account in the group recommendation algorithm helps to rule out items which group members consider unacceptable (e.g., in the context of music recommendation [8]).

Preferences in Constraint-Based Recommendation This type of recommender system is used in situations where items and recommendation knowledge is specified on a semantic level, for example, in terms of rules. In single-user as well as in group settings, preferences can be specified on the level of item attributes or user requirements that are related to item properties. In most of the cases, such preferences are represented in terms of specific types of rules [21]. Preferences between item attributes can also be specified on the basis of preference networks [7]. Attribute weights and interest dimensions are often used in the context of a utility-based analysis of recommendation candidates derived from a constraint-based recommendation process [22]. Preference collection in group-based recommendation settings resembles single-user settings, however, mechanisms are needed to resolve inconsistencies between the preferences of group members (see also Chap. 2).

Preferences in Critiquing-Based Recommendation Critiques are collected to derive user-individual recommendations. These can be aggregated afterwards to build a group model that is used for determining group recommendations [51]. Critiques can be specified directly on item attributes via conventional mechanisms such as *compound critiques* or *unit critiques* or on the basis of more advanced concepts such as *natural language based critiques* [31]. Such types of critiques can also be used in group recommendation. Natural language interfaces for group decision support have not been investigated up to now. Further information that can be used to understand preferences is provided in chat-based approaches [55].

5.3 Preference Handling Practices

Types of Preferences Ratings are influenced by the current context of the user [3, 6]. Some examples of contextual factors are (1) the time between item consumption and item evaluation (the longer the time, the more ratings regress towards the middle of the scale [6]) and (2) the type of rating scale used. Anchoring biases (see Chap. 8) can, for example, be reduced by applying binary or star-based rating scales (compared to numerical rating scales [1]). In general, adapted rating scales and preference collection user interfaces help to avoid rating biases, compared to post-hoc debiasing algorithms [1]. An analysis of anchoring effects based on rating interface is also presented in Cosley et al. [14]. The authors show that item evaluations by other users have an impact on the rating behavior of the current user (if made visible). The existence of the *positivity effect* in the recommendation context, i.e., pleasant items are processed and recalled from memory more effectively, is shown in [6]. In the context of group recommender systems, it has also been shown that multi-attribute utility-based rating scales can help to make ratings more stable in terms of a lower standard deviation of individual evaluations [38, 70]. In the context of critiquing-based recommender systems, combined preference feedback such as compound critiques and natural language based feedback helps to significantly reduce the number of critiquing cycles needed by a user to find a relevant item [31, 50]. In conversational recommendation scenarios [13], users specify their preferences in terms of preferred attribute values. In this context, not all attributes are of relevance for each user. For example, in a digital camera recommender, a user might be interested in specifying the desired camera type and resolution but not in specifying the supported video formats (reasons could include the irrelevance of video functionalities for his/her work, or a limited amount of technical domain knowledge). Approaches to recommending which questions/parameters to be shown to users are presented in [20, 21, 45]. Finally, in content-based recommendation, additional knowledge about user preferences collected, for example, in the form of eye-tracking data, can help to significantly improve the prediction quality of the recommendation algorithm [77].

Visibility of Preferences In the context of group decision making, we face the question of how to disclose the preferences of individual group members to other group members [36, 69]. Group members could be interested in seeing the preferences of other group members for different reasons. For example, if there are some experts in the group, non-experts engaged in the decision making process would like to follow the experts (effort-saving aspect [36]). Furthermore, what a single group member wants can depend directly on what other group members want. For example, if one group member likes to play tennis, his/her interest in having a hotel that offers a tennis court depends on the existence of other group members interested in playing tennis. If no other group members are interested in tennis, preferences regarding having a tennis court become moot. However, the other side of the coin is that knowing the preferences of other group members can lead to situations where

potentially decision-relevant knowledge is not made available to all group members due to a focus shift towards analyzing the preferences of other group members [69]. Furthermore, if some group members are able to communicate negative feedback to all group members, phenomena such as *emotional contagion* [49] can occur, i.e., other group members can be infected by negative moods. There is also the danger of *GroupThink* by which strongly coherent groups try to avoid conflicts and therefore agree on already established preferences. As a consequence, preference visibility should be postponed until individual group members have articulated their own preferences with regard to a set of items [69]. Following this approach, the *overall satisfaction* with the outcome of a group decision process can be increased and *anchoring effects* can be reduced, since group members focus more on item evaluation than on the analysis of the preferences of other group members [69]. Postponed preference visibility in collaborative preference specification processes also leads to an increased exchange of decision-relevant knowledge which helps to improve the overall *quality of a decision* [4]. An additional factor to increase the amount of content/knowledge exchange is *recommendation diversity*. In the extreme case, when recommendations reflect opinions that completely contradict the currently-defined preferences of group members, the amount of information exchanged between group members reaches its maximum [28]. How much diversity is accepted by a user (or a group), is still an open issue for future research.

Choice Overload The basic idea underlying the notion of choice overload is that the higher the number of decision alternatives (i.e., items shown by a collaborative and content-based recommender or parameters shown by a constraint-based or critiquing-based recommender), the higher the related effort to analyze alternatives, and the lower the probability that a decision is made (due to *choice overload*) [17, 34, 66]. Bollen et al. [6] analyzed the role of choice overload in the context of collaborative filtering based recommendation scenarios. They detected that larger result sets containing only attractive items do not necessarily lead to higher choice satisfaction compared to smaller item sets. In other words, the increasing attractiveness of result sets is counteracted by an increase in effort. The authors mention an optimal result set size of 5–7 but explicitly point out the need for further related research. A meta-analysis on choice overload [66] showed that choice overload is not omnipresent and that it mostly occurs when alternatives are very similar and users lack sufficient expertise to have stable and clear preferences. Later work by Willemsen et al. [74] showed that latent feature diversification can reduce choice difficulty and improve satisfaction. The diversification method reduced the similarity between items while controlling for their attractiveness, making small sets just as attractive and satisfactory as larger sets, with much less choice difficulty. For group decisions, choice overload could be tackled in creative ways, extending the diversification methods used for single users. One could imagine, for example, giving each group member a (diverse) subset of items out of which the best items should be identified. Afterwards, the group as a whole can decide which options to select from the conjunction of the best items from each of the subsets. In this way, resources of individual decision makers are combined, and larger sets of items can

be handled without much risk of choice overload. Consequently, reducing choice overload by using the joint resources of a group is an interesting new research direction.

In the context of constraint-based and critiquing-based recommender systems, similar studies are needed focusing on aspects such as result set size, but also on number of questions posed to the user and number of different repair alternatives shown in situations where no solution can be found by the recommender system. Mechanisms to reduce the number of questions are presented in [20, 21, 45] where questions are selected on the basis of collaborative recommendation algorithms [20, 21], or where information-gain based measures are used to predict the next relevant questions to be posed to users [45]. Groups often apply choice deferral more frequently than individuals [73]. As mentioned in White et al. [73], possible explanations thereof are (1) defending a choice deferral seems to be easier and easier to justify than the selection of an option. For example, in jury decision making there is often a tendency towards acquittal. (2) Groups are more risk-seeking than individuals (see Chap. 8), and choice deferral is often a riskier behavior. (3) Groups as a whole often have more reasons to defer a decision compared to individuals. Finally, we want to point out that the optimal size of a choice set can also differ, depending on item selection strategy. For example, users who emphasize finding the optimal solution (maximizers) would like to analyze *as many items as possible* whereas users interested in finding a satisfying solution as quickly as possible (satisficers) prefer smaller option sets [67].

5.4 Consistency Management

There exist situations where no solution / recommendation can be found for a given set of user requirements, especially in the context of constraint-based recommendation scenarios [21]. Given, for example, a set of user requirements (represented by a list of attribute/value pairs) which is *inconsistent* with the underlying product catalog (e.g., pre-defined item list), a user needs support to know which attribute values have to be adapted in order to be able to identify a solution [20]. In such scenarios, conflict detection and diagnosis techniques can help to automatically figure out minimal sets of requirements that have to be adapted in order to find a solution [24, 27, 63]. Whereas [24, 27] focus on the determination of personalized diagnoses for single users, [29] introduced an approach to take into account the principles of computational social choice [12] for diagnosing inconsistent user requirements in group-based recommendation and configuration settings (for example, diagnosis ranking is implemented on the basis of *least misery*). In group-based settings, inconsistencies do not only occur between user requirements and the underlying product catalog, but also between the requirements / preferences of different group members [29]. Similar inconsistencies can occur in critiquing-based recommendation scenarios. For example, if the complete critiquing history of a user (or a group [51]) is used to calculate recommendations, inconsistencies between

critiques have to be resolved. In most cases, such inconsistencies are resolved by simply omitting elder critiques and leaving the more recent ones in the active set. Diagnoses for inconsistent requirements can also be regarded as an explanation that can help users out of the *no solution could be found* dilemma [24]. Such explanations can help to make the identification of relevant items more efficient and can also help to increase the *trust* of a user and the degree of *domain knowledge*, which is extremely important in order to make high-quality decisions [23].

5.5 Conclusions and Research Issues

In this chapter, we focused on a short overview of existing approaches to support the handling of preferences. Preference handling mechanisms from single-user recommendation scenarios can often be applied in group-based settings, but more work is needed to investigate how preference elicitation procedures can be optimized for the group recommendation context. Furthermore, we summarized insights from user studies focusing on the acquisition of preferences and also on the management of inconsistent user requirements, i.e., requirements for which a recommender cannot find a solution. In this context, there are a couple of open research issues which will be discussed in the following.

There exist a couple of research contributions that introduce and discuss *aggregation mechanisms* that can be used to integrate individual user preferences. For example, in [47, 48] Masthoff introduces social-choice-based aggregation mechanisms (e.g., *Least Misery* (LMS)—see Chap. 2) that can be used to identify recommendations for a group. Although initial insights have already been provided in terms of which aggregation mechanisms are useful [48], there is no in-depth analysis of which aggregation strategies should be applied in which context. An analysis of the appropriateness of aggregation strategies depending on item type is presented in Felfernig et al. [30]. A related insight is that, for decisions related to high-involvement items, groups tend to apply *Least Misery*-style heuristics, whereas in low-involvement item domains, misery of individual users is accepted to a larger extent. Two examples of aggregation methods used in this context are *Average* (AVG) and *Most Pleasure* (MPL). An open issue in this context is how to integrate basic aggregation functions with knowledge of the personality and emotions of group members (see also [62]). New related insights will serve as a basis for context-dependent preference aggregation mechanisms that take into account the group context before deciding which aggregation and corresponding explanation method to apply.

Avoiding manipulations is an important aspect of assuring high-quality, fair group decision making. In order to achieve this goal, aggregation mechanisms have to be provided (in combination with corresponding recommender user interfaces) that help to avoid different kinds of manipulation efforts. Related work in the context of group recommendation has already been performed by Jameson et al. [35]. For example, median-based aggregation heuristics help to avoid an impact

of extremely high or low item evaluations. Further mechanisms can be included to limit the number of possible item evaluations per group member and to give feedback on the current status of the decision process on a meta-level. For example, in terms of statements such as *user X changed his/her preferences N times with regard to item A, the evaluations range from 1 to 4 stars*. A question that has to be answered in this context is to which extent we have to adapt user interfaces from single user recommendation scenarios to the group context [49]. For example, in which context should one provide information about the preferences of other group members or information about specific inconsistencies between the preferences of group members. Although user interfaces provide different mechanisms to handle user and group preferences, additional approaches have to be developed to improve the quality of the group decision making processes. For example, approaches that *better predict the preferences of the group, improve the quality of the decision outcome, and enable a more efficient process towards the achievement of group consensus*. User interfaces should also be capable of stimulating intended behavior, for example, stimulating information exchange between group members in order to make decision-relevant knowledge available to the whole group [28].

References

1. G. Adomavicius, J. Bockstedt, C. Shawn, J. Zhang, De-biasing user preference ratings in recommender systems, in *Joint Workshop on Interfaces and Human Decision Making for Recommender Systems*. CEUR Workshop Proceedings, vol. 1253 (2014), pp. 2–9
2. M. Agarwal, D. Reid, Predicting group choice: an experimental study using conjoint analysis, in *Academy of Marketing Science (AMS) Annual Conference* (1984), pp. 445–449
3. X. Amatriain, J. Pujol, N. Tintarev, N. Oliver, Rate it again: increasing recommendation accuracy by user re-rating, in *3rd ACM Conference on Recommender Systems*, New York, USA (2009), pp. 173–180
4. M. Atas, A. Felfernig, M. Stettinger, T.N. Trang Tran, Beyond item recommendation: using recommendations to stimulate knowledge sharing in group decisions, in *9th International Conference on Social Informatics (SocInfo 2017)*, Oxford, UK (2017), pp. 368–377
5. J. Bettman, M. Luce, J. Payne, Constructive consumer choice processes. *J. Consum. Res.* **25**(3), 187–217 (1998)
6. D. Bollen, M. Graus, M. Willemsen, Remembering the stars?: effect of time on preference retrieval from memory, in *6th CM Conference on Recommender Systems* (Dublin, Ireland, 2012), pp. 217–220
7. R. Brafman, F. Rossi, D. Salvagnin, K. Venable, T. Walsh, Finding the next solution in constraint- and preference-based knowledge representation formalisms, in *12th International Conference on the Principles of Knowledge Representation and Reasoning (KR 2010)*, Toronto, Ontario, Canada (2010), pp. 425–433
8. D. Chao, J. Balthorp, S. Forrest, Adaptive radio: achieving consensus using negative preferences, in *ACM SIGGROUP Conference on Supporting Group Work*, Sanibel Island, FL, USA (2005), pp. 120–123
9. L. Chen, P. Pu, Critiquing-based recommenders: survey and emerging trends. *User Model. User-Adap. Inter.* **22**(1–2), 125–150 (2012)
10. L. Chen, G. Chen, F. Wang, Recommender systems based on user reviews: the state of the art. *User Model. User-Adap. Inter.* **25**(2), 99–154 (2015)

11. L. Chen, F. Wang, W. Wu, Inferring users' critiquing feedback on recommendations from eye movements, in *International Conference on Case-Based Reasoning (ICCBR'16)*, Atlanta, GA, USA (2016), pp. 62–76
12. Y. Chevalere, U. Endriss, J. Lang, N. Maudet, A short introduction to computational social choice, in *33rd Conference on Current Trends in Theory and Practice of Computer Science*, Harrachov, Czech Republic (2007), pp. 51–69
13. K. Christakopoulou, F. Radlinski, K. Hofmann, Towards conversational recommender systems, in *International Conference on Knowledge Discovery and Data Mining (KDD 2016)*, San Francisco, CA, USA (2016), pp. 815–824
14. D. Cosley, S. Lam, I. Albert, J. Konstan, J. Riedl, Is seeing believing?: how recommender system interfaces affect users' opinions, in *CHI'03* (2003), pp. 585–592
15. A. Crossen, J. Budzik, K. Hammond, FLYTRAP: intelligent group music recommendation, in *7th International Conference on Intelligent User Interfaces*, San Francisco, CA, USA (2002), pp. 184–185
16. M. De Gemmis, L. Iaquinta, P. Lops, C. Musto, F. Narducci, G. Semeraro, Preference learning in recommender systems, in *ECML/PKDD-09 Workshop* (2009), pp. 41–55
17. K. Diehl, C. Poynor, Great expectations?! assortment size, expectations, and satisfaction. *J. Mark. Res.* **47**(2), 312–322 (2010)
18. M.D. Ekstrand, M. Willemsen, Behaviorism is not enough: better recommendations through listening to users, in *Proceedings of the 10th ACM Conference on Recommender Systems, RecSys '16*, Boston, MA, USA (ACM, New York, 2016), pp. 221–224
19. M. Ekstrand, J. Riedl, J. Konstan, Collaborative filtering recommender systems. *Found. Trends Hum. Comput. Interact.* **4**(2), 81–173 (2011)
20. A. Falkner, A. Felfernig, A. Haag, Recommendation technologies for configurable products. *AI Mag.* **32**(3), 99–108 (2011)
21. A. Felfernig, R. Burke, Constraint-based recommender systems: technologies and research issues, in *ACM International Conference on Electronic Commerce (ICEC08)*, Innsbruck, Austria (2008), pp. 17–26
22. A. Felfernig, G. Friedrich, D. Jannach, M. Zanker, An integrated environment for the development of knowledge-based recommender applications. *Int. J. Electron. Commer.* **11**(2), 11–34 (2006)
23. A. Felfernig, B. Gula, E. Teppan, Knowledge-based recommender technologies for marketing and sales. *Int. J. Pattern Recognit. Artif. Intell.* **21**(2), 1–22 (2006). Special Issue of Personalization Techniques for Recommender Systems and Intelligent User Interfaces
24. A. Felfernig, G. Friedrich, M. Schubert, M. Mandl, M. Mairitsch, E. Teppan, Plausible repairs for inconsistent requirements, in *IJCAI'09*, Pasadena, CA (2009), pp. 791–796
25. A. Felfernig, M. Schubert, C. Zehentner, An efficient diagnosis algorithm for inconsistent constraint sets. *Artif. Intell. Eng. Des. Anal. Manuf.* **26**(1), 53–62 (2012)
26. A. Felfernig, M. Jeran, G. Ninaus, F. Reinfrank, S. Reiterer, M. Stettinger, Basic approaches in recommendation systems. *Recommendation Systems in Software Engineering* (Springer, Berlin, 2013), pp. 15–37
27. A. Felfernig, M. Schubert, S. Reiterer, Personalized diagnosis for over-constrained problems, in *23rd International Conference on Artificial Intelligence (IJCAI 2013)*, Peking, China (2013), pp. 1990–1996
28. A. Felfernig, M. Stettinger, G. Leitner, Fostering knowledge exchange using group recommendations, in *ACM RecSys'15 Workshop on Interfaces and Human Decision Making for Recommender Systems (InRS'15)*, Vienna, Austria (2015), pp. 9–12
29. A. Felfernig, M. Atas, T.N. Trang Tran, M. Stettinger, Towards group-based configuration, in *International Workshop on Configuration 2016 (ConfWS'16)* (2016), pp. 69–72
30. A. Felfernig, M. Atas, T.N. Trang Tran, M. Stettinger, S. Polat-Erdeniz, An analysis of group recommendation heuristics for high- and low-involvement items, in *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems (IEA/AIE 2017)*, Arras, France (2017), pp. 335–344

31. P. Grasch, A. Felfernig, F. Reinfrank, ReComment: towards critiquing-based recommendation with speech interaction, in *7th ACM Conference on Recommender Systems* (ACM, New York, 2013), pp. 157–164
32. M. Graus, M. Willemsen, Improving the user experience during cold start through choice-based preference elicitation, in *Proceedings of the 9th ACM Conference on Recommender Systems, RecSys '15*, Vienna, Austria (ACM, New York, 2015), pp. 273–276
33. D. Grether, C. Plott, Economic theory of choice and the preference reversal phenomenon. *Am. Econ. Rev.* **69**(4), 623–638 (1979)
34. C. Huffman, B. Kahn, Variety for sale: mass customization or mass confusion? *J. Retail.* **74**(4), 491–513 (1998)
35. A. Jameson, More than the sum of its members: challenges for group recommender systems, in *International Working Conference on Advanced Visual Interfaces* (2004), pp. 48–54
36. A. Jameson, B. Smyth, Recommendation to groups, in *The Adaptive Web*, ed. by P. Brusilovsky, A. Kobsa, W. Nejdl. Lecture Notes in Computer Science, vol. 4321 (Springer, New York, 2007), pp. 596–627
37. A. Jameson, S. Baldes, T. Kleinbauer, Two methods for enhancing mutual awareness in a group recommender system, in *ACM International Working Conference on Advanced Visual Interfaces*, Gallipoli, Italy (2004), pp. 447–449
38. A. Jameson, M. Willemsen, A. Felfernig, M. de Gemmis, P. Lops, G. Semeraro, L. Chen, Human decision making and recommender systems, in *Recommender Systems Handbook*, 2nd edn., ed. by F. Ricci, L. Rokach, B. Shapira. (Springer, Berlin, 2015), pp. 611–648
39. D. Jannach, M. Zanker, A. Felfernig, G. Friedrich, *Recommender Systems – An Introduction* (Cambridge University Press, Cambridge, 2010)
40. G. Jawaheer, P. Weller, P. Kostkova, Modeling user preferences in recommender systems: a classification framework for explicit and implicit user feedback. *ACM Trans. Interact. Intell. Syst.* **4**(2), 1–26 (2014)
41. S. Kalloori, F. Ricci, M. Tkalcić, Pairwise preferences based matrix factorization and nearest neighbor recommendation techniques, in *Proceedings of the 10th ACM Conference on Recommender Systems, RecSys '16*, New York, NY, USA (ACM, New York, 2016), pp. 143–146
42. D. Kluver, T. Nguyen, M. Ekstrand, S. Sen, J. Riedl, How many bits per rating?, in *Proceedings of the Sixth ACM Conference on Recommender Systems, RecSys '12*, New York, NY, USA (ACM, New York, 2012), pp. 99–106
43. E. Kurdyukova, S. Hammer, E. André, Personalization of content on public displays driven by the recognition of group context. *Ambient Intell.* **7683**, 272–287 (2012)
44. S. Lichtenstein, P. Slovic, *The Construction of Preference* (Cambridge University Press, Cambridge, 2006)
45. T. Mahmood, F. Ricci, Improving recommender systems with adaptive conversational strategies, in *20th ACM Conference on Hypertext and Hypermedia*, Torino, Italy (2009), pp. 73–82
46. J. Masthoff, Modeling the multiple people that are me, in *User Modeling 2003*. Lecture Notes in Artificial Intelligence, vol. 2702 (Springer, Berlin, 2003), pp. 258–262
47. J. Masthoff, Group recommender systems: combining individual models, in *Recommender Systems Handbook* (Springer, London, 2011), pp. 677–702
48. J. Masthoff, Group recommender systems: aggregation, satisfaction and group attributes, in *Recommender Systems Handbook* (Springer, London, 2015), pp. 743–776
49. J. Masthoff, A. Gatt, In pursuit of satisfaction and the prevention of embarrassment: affective state in group recommender systems. *User Model. User-Adap. Inter.* **16**(3–4), 281–319 (2006)
50. K. McCarthy, J. Reilly, L. McGinty, B. Smyth, On the dynamic generation of compound critiques in conversational recommender systems, in *International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems* (Springer, Berlin, 2004), pp. 176–184
51. K. McCarthy, L. McGinty, B. Smyth, M. Salamó, Social interaction in the CATS group recommender, in *Workshop on the Social Navigation and Community based Adaptation Technologies* (2006)

52. D. McFadden, Rationality for economists. *J. Risk Uncertain.* **19**(1–3), 73–105 (1999)
53. B. Miller, I. Albert, S. Lam, J. Konstan, J. Riedl, MovieLens unplugged: experiences with a recommender system on four mobile devices, in *People and Computers XVII - Designing for Society*, ed. by E. O'Neill, P. Palanque, P. Johnson (Springer, London, 2004), pp. 263–279
54. J. Neidhardt, L. Seyfang, R. Schuster, H. Werthner, A picture-based approach to recommender systems. *Inform. Technol. Tour.* **15**(1), 49–69 (2015)
55. T. Nguyen, F. Ricci, A chat-based group recommender system for tourism, in *Information and Communication Technologies in Tourism*, ed. by R. Schegg, B. Stangl (Springer, Cham, 2017), pp. 17–30
56. J. Payne, J. Bettman, E. Johnson, *The Adaptive Decision Maker* (Cambridge University Press, Cambridge, 1993)
57. M. Pazzani, D. Billsus, Learning and revising user profiles: the identification of interesting web sites. *Mach. Learn.* **27**(3), 313–331 (1997)
58. B. Peintner, P. Viappiani, N. Yorke-Smith, Preferences in interactive systems: technical challenges and case studies. *AI Mag.* **29**(4), 13–24 (2008)
59. A. Pommeranz, J. Broekens, P. Wiggers, W. Brinkman, C. Jonker, Designing interfaces for explicit preference elicitation: a user-centered investigation of preference representation and elicitation process. *User Model. User-Adap. Inter.* **22**(4–5), 357–397 (2012)
60. P. Pu, L. Chen, User-involved preference elicitation for product search and recommender systems. *AI Mag.* **29**(4), 93–103 (2008)
61. Z. Qiao, P. Zhang, Y. Cao, C. Zhou, L. Guo, Improving collaborative filtering recommendation via location-based user-item subgroup. *Proc. Comput. Sci.* **29**, 400–409 (2014)
62. L. Quijano-Sanchez, J. Recio-García, B. Díaz-Agudo, G. Jiménez-Díaz, Social factors in group recommender systems. *ACM Trans. Intell. Syst. Technol.* **4**(1), 8:1–8:30 (2006)
63. R. Reiter, A theory of diagnosis from first principles. *AI J.* **32**(1), 57–95 (1987)
64. F. Ricci, Q. Nguyen, Acquiring and revising preferences in a critique-based mobile recommender systems. *IEEE Intell. Syst.* **22**(3), 22–29 (2007)
65. P. Sawyer, S. Viller, I. Sommerville, A behavioral model of choice. *Q. J. Econ.* **69**(1), 99–118 (1955)
66. B. Scheibehenne, R. Greifeneder, P. Todd, Can there ever be too many options? a meta-analytic review of choice overload. *J. Consum. Res.* **37**(3), 409–425 (2010)
67. B. Schwartz, A. Ward, J. Monterosso, S. Lyubomirsky, K. White, D. Lehman, Maximizing versus satisficing: happiness is a matter of choice. *J. Pers. Soc. Psychol.* **83**(5), 1178–1197 (2002)
68. E. Sparling, S. Sen, Rating: How difficult is it?, in *5th ACM Conference on Recommender Systems* (ACM, New York, 2011), pp. 149–156
69. M. Stettinger, A. Felfernig, G. Leitner, S. Reiterer, Counteracting anchoring effects in group decision making, in *23rd Conference on User Modeling, Adaptation, and Personalization (UMAP'15)*. Lecture Notes in Computer Science, vol. 9146 (Dublin, Ireland, 2015), pp. 118–130
70. M. Stettinger, A. Felfernig, G. Leitner, S. Reiterer, M. Jeran, Counteracting serial position effects in the CHOICLA Group decision support environment, in *20th ACM Conference on Intelligent User Interfaces (IUI2015)*, Atlanta, Georgia, USA (2015), pp. 148–157
71. T. Ulz, M. Schwarz, A. Felfernig, S. Haas, A. Shehadeh, S. Reiterer, M. Stettinger, Human computation for constraint-based recommenders. *J. Intell. Inf. Syst.* **49**(1), 37–57 (2017)
72. S. Wei, N. Ye, Q. Zhang, Time-aware collaborative filtering for recommender systems, in *Chinese Conference on Pattern Recognition*, Beijing, China (2012), pp. 663–670
73. C. White, S. Hafenbrädl, U. Hoffrage, N. Reisen, J. Woike, Are groups more likely to defer choice than their members. *Judgm. Decis. Mak.* **6**(3), 239–251 (2011)
74. M. Willemsen, M. Graus, B. Knijnenburg, Understanding the role of latent feature diversification on choice difficulty and satisfaction. *User Model. User-Adap. Inter.* **26**(4), 347–389 (2016)

75. D. Winterfeldt, W. Edwards, *Decision Analysis and Behavioral Research* (Cambridge University Press, Cambridge, 1986)
76. H. Xie, J. Lui, Mathematical modeling of group product recommendation with partial information: How many ratings do we need? *Perform. Eval.* **77**, 72–95 (2014)
77. S. Xu, H. Jiang, F. Lau, Personalized online document, image and video recommendation via commodity eye-tracking, in *ACM Conference on Recommender Systems (RecSys'08)*, Lausanne, Switzerland (2008), pp. 83–90