

Chapter 6

Explanations for Groups

Alexander Felfernig, Nava Tintarev,
Thi Ngoc Trang Tran, and Martin Stettinger^{ab}

^a Citation: Alexander Felfernig, Nava Tintarev, Thi Ngoc Trang Tran, and Martin Stettinger. Explanations for Groups, in: Group Recommender Systems – An Introduction, Alexander Felfernig, Ludovico Boratto, Martin Stettinger, and Marko Tkalčić (eds.), Springer, pp. 105–126, ISBN: 978-3-319-75066-8, 2018.

^b This is a pre-print version of the chapter published in the book "Group Recommender Systems: An Introduction": <http://www.springer.com/us/book/9783319750668>.

Abstract Explanations are used in recommender systems for various reasons. Users have to be supported in making (high-quality) decisions more quickly. Developers of recommender systems want to convince users to purchase specific items. Users should better understand how the recommender system works and why a specific item has been recommended. Users should also develop a more in-depth understanding of the item domain. Consequently, explanations are designed in order to achieve specific *goals* such as increasing the transparency of a recommendation or increasing a user's trust in the recommender system. In this chapter, we provide an overview of existing research related to explanations in recommender systems, and specifically discuss aspects relevant to group recommendation scenarios. In this context, we present different ways of explaining and visualizing recommendations determined on the basis of *aggregated predictions* and *aggregated models* strategies.

6.1 Introduction

Explanations have been recognized as an important means to help users to evaluate recommendations, and make better decisions, but also to deliver persuasive messages to the user [30, 62]. Empirical studies show that users appreciate explanations of recommendations [14, 30]. Explanations can be regarded as *a means to make something clear by giving a detailed description* [63]. In the recommender systems context, Friedrich and Zanker [26] define explanations as *information about recommendations* and as *means to support objectives defined by the designer of a recommender system*. Explanations can be seen from two basic viewpoints [5, 65]: (1) the *user's (group member's)* and (2) the *recommender provider's* point of view. *Users of recommender systems* are in the need of additional information to be able to develop a better understanding of the recommended items. *Developers of recommender systems* want to provide additional information to users for various reasons,

for example, to convince the user to purchase an item, to increase a user's item domain knowledge (educational aspect), and to increase a user's *trust in* and overall *satisfaction with* the recommender system. Another objective is to make users more *tolerant* with regard to recommendations provided by the system. This is especially important for new users/items, otherwise a recommendation may be perceived as inappropriate. Solely providing the core functionality of recommender systems, i.e., showing a list of relevant items to users, could evoke the impression of interacting with a *black box* with no transparency and no additional user-relevant information [30, 62]. Consequently, explanations are an important means to provide information related to recommendations, the recommendation process, and further objectives defined by the designer of a recommender system [12, 26, 38, 53, 67]. Visualizations of explanations can further improve the perceived quality of a recommender system [27, 65, 67] – where appropriate, examples of visualizations will be provided.

Explanations in Single User Recommender Systems

In single user recommender systems, various efforts have already been undertaken to categorize explanations with regard to *information sources used to generate explanations* and corresponding *goals of explanations* [26, 28, 48, 61, 62, 64]. A categorization of different information sources that can be used for the explanation of recommendations is given, for example, in Friedrich and Zanker [26] where *recommended items*, *alternative items*, and the *user model* are mentioned as three orthogonal *information categories*. Potential goals of explanations are discussed a.o. in Tintarev and Masthoff [62] and Jameson et al. [34]. Examples thereof are *efficiency* (reducing the time needed to complete a choice task), *persuasiveness* (exploiting explanations to change a user's choice behavior) [29], *effectiveness* (proactively helping the user to make higher-quality decisions), *transparency* (reasons as to why an item has been recommended, i.e., answering why-questions), *trust* (supporting a user in increasing her confidence in the recommender), *scrutability* (providing ways to make the user profile manageable), *satisfaction* (explanations focusing on aspects such as enjoyment and usability), and *credibility* (assessed likelihood that a recommendation is accurate). Bilgic and Mooney [5] offer a differentiation between explanations that focus on (1) *promotion*, i.e., convincing users to adopt recommendations, and (2) *satisfaction*, i.e., to help users make more accurate decisions.

Examples of verbal explanations for single user recommendations include phrases such as (1) '*users who purchased item x also purchased item y*', (2) '*since you liked the book x, we recommend book y from the same authors*', (3) '*since you prefer taking sports photos, we recommend camera y because it supports 10 pics/sec in full-frame resolution*', and (4) '*item y would be a good choice since it is similar to the already presented item x and has the requested higher frame rate (pics/sec)*'. These example explanations are formulated based on information collected and provided by the underlying recommendation approaches, i.e., (1) collaborative filtering, (2) content-based filtering, (3) constraint-based recommendation, and (4) critiquing-based recommendation – see, for example, [11, 19, 28, 30]. These examples of ex-

planations can be regarded as 'basic', since further information could be included. For instance, information related to competitor items and previous user purchases: *'since you prefer taking sports photos, we recommend camera y because it supports 10 pics/sec in full-frame resolution. z would have been the other option but we propose y since you preferred purchasing from provider k in the past and y is only a little bit more expensive than its competitors'*.

Another type of explanation is the following: *'no solution could be found – if you increase the maximum acceptable price or decrease the minimum acceptable resolution, a corresponding solution can be identified.'* This explanation focuses on indicating options to find a way out of the 'no solution could be found' dilemma which primarily occurs in the context of constraint-based recommendation scenarios [18]. Another example is *'item y outperforms item z in both, quality and price, whereas x outperforms z only in quality'*. This explanation does not focus on one item but supports the *comparison* of different candidate items (in this case, *x* and *y*). Importantly, it is directly related to the concept of *asymmetric dominance* (*y* outperforms *z* two times whereas *x* does this only once) which is a *decision bias* discussed in Chapter 8. Explanations based on *item comparisons* are mostly supported in critiquing-based [11] and constraint-based recommendation [21] which are both based on semantic recommendation knowledge (see Chapters 1 and 2). In critiquing-based recommendation, *compound critiques* point out the relationship between the current reference item and the corresponding candidate items [43]. An example of a compound critique in the domain of *digital cameras* is the following: *on the basis of the current reference item x, you can take a look at cameras with a [lower price] and a [higher resolution] or at cameras with a [higher price] and a [higher optical zoom]*. An analysis of comparison interfaces in single user constraint-based recommendation is presented in [21, 22].

Explanations in Group Recommender Systems

The aforementioned explanation approaches focus on single users, and so, do not have to consider certain aspects of group decision making. Explanations for groups can have *further goals* such as *fairness* (taking into account, as far as possible, the preferences of all group members), *consensus* (group members agree on the decision), and *optimality* (a group makes an optimal or nearly-optimal decision¹). An important aspect in this context is that explanations show how the interests of individual group members are taken into account. This is not relevant in the context of single user recommender systems. Understanding the underlying process enables group members to evaluate the appropriateness of the way their preferences have to been taken into account by the group recommender system. Similar to explanations for single users, explanations for groups are shaped by the underlying recommendation algorithms. Explanations similar to those already mentioned can also be defined in a group context. For example, (1) *'groups that like item x also like item y'*, (2)

¹ In contrast to single-user decision making, the exchange of decision-relevant knowledge among group members has to be fostered [4].

'since the group likes the film x , we also recommend film y from the same director', (3) 'since the maximum camera price accepted by group members is 500 (defined by Paul) and the minimum accepted resolution is 18 mpix (defined by Joe), we recommend y which supports 20 mpix at a price of 459.', and (4) 'item x is a good choice since it supports a higher frame rate requested by all group members and is only a little bit more expensive'.

These examples show that the chosen preference aggregation approach (see Chapter 2) has an impact on the explanation style. While *aggregated predictions* include information about the individual preferences of group members (e.g., one group member specified the lowest maximum price of 500) and thus support explanation goals such as *fairness* and *consensus*, *aggregated models*-based approaches restrict explanations to the group level (e.g., groups that like x also like y). More advanced (hybrid) explanations [37] can also be formulated in group recommendation scenarios, for example, 'since all group members prefer sports photography, we recommend camera y rather than camera z . It is only a little bit more expensive but has a higher usability which is important for group member Joe who is a newbie in digital photography. Similar groups also preferred y '.

An example of an explanation in a situation where no solution could be found is: 'no 23 mpix camera with a price below 250 could be found. Therefore we recommend camera y with 20 mpix and a price of 249 since price is the most important criterion for all group members.' Finally, the following example shows how to take into account a group's social reality, for example, in terms of 'tactful' explanations [53]: 'Although your preference for item y is not very high, your close friend Peter thinks it is an excellent choice'. This example explanation is formulated on the level of *aggregated predictions* (see Chapter 2) and also takes into account social relationships among group members (e.g., neighborhoods in a social network). On the level of *aggregated models*, an explanation can be formulated as follows: 'A majority thinks that it is a good choice. Some group members think that it is an excellent choice.' (assuming the existence of at least some aggregated categorization of preferences such as *number of likes*). Taking into account the individual preferences of group members helps to increase *mutual awareness* among group members, and thus counteracts the natural tendency to focus on one's own favorite alternatives [33]. An approach to explaining the *consequences of a given recommendation* is introduced by Jameson et al. [32], where *emotions* of individual group members with regard to a recommendation are visualized in terms of animated characters.

We want to emphasize that *explanations for groups* is a highly relevant research topic with a limited, but nevertheless direction-giving, number of research results [3, 13, 31, 33, 47]. In the following, we sketch ways in which explanations for single-user recommendation scenarios can be adapted to groups. Following the idea of categorizing explanation types along the different recommendation approaches [63, 68], we discuss explanations for groups in the context of *collaborative*- and *content-based filtering*, as well as *constraint*- and *critiquing-based recommendation*.

6.2 Collaborative Filtering

A widely used example of explanations in collaborative filtering recommenders is *'users who purchased item x also purchased item y'*. Such explanations can be generated, for example, on the basis of *association rule mining* which is often used as a model-based collaborative filtering approach [40]. Herlocker et al. [30] analyzed the role of explanations in collaborative filtering recommenders. They focused on the impact of different explanation styles on user acceptance of recommender systems. Explanations were mostly represented graphically. For example, a histogram of neighbors' ratings for the recommended item categorized ratings as 'good', 'neutral', or 'bad'. The outcome of their study was that rating histograms are the most compelling way to explain rating data. Furthermore, *simple graphs* were perceived as more compelling than more detailed explanations, i.e., *simplicity of explanations is a key factor*.

An orthogonal approach to propose explanations for collaborative-filtering-based recommendations is presented by Chang et al. [10]. Following the idea of generating recommendations based on knowledge from the crowd (see, e.g., [66]), the authors introduce the idea of asking crowd workers to provide feedback on explanations. Quality assurance is an issue but crowd-sourced explanations were considered high-quality. The authors mention *longer explanation texts* and an *increased number of references to item genres* as examples of indicators of high-quality explanations. An example of a question for crowd-sourcing in group recommendation scenarios is the following: *'given this movie recommendation (e.g., Guardians of the Galaxy), which of the following are useful explanations for a group of middle-aged persons? Can be viewed by the whole family; Includes plenty of songs from the 70ies; Best movie we have ever seen'*. This way, crowd knowledge can be exploited to better figure out which kinds of explanations are useful in which context and which ones might be particularly well-received by specific groups (in this case, middle-aged persons). A similar approach can be used to figure out relevant explanations in other recommendation approaches, i.e., *which tags to use for an explanation?* (content-based filtering), *which requirements to relax?* (constraint-based recommendation), and *which critiques to propose to the user?* (critiquing-based recommendation).

As mentioned by Bilgic and Mooney [5], a goal of the explanations introduced in Herlocker et al. [30] is to promote items but not to provide more insights as to why the items have been recommended, i.e., not to provide satisfaction-oriented explanations that might help users to make more accurate decisions. There are different ways to move the explanation focus towards more informative explanations. As proposed in [5] (for single user recommenders), a collaborative-filtering-based explanation can be extended by providing information on items that had a major influence on the determination of the proposed recommendation. Removing the *most influential items* (already rated by group members) from the set of rated items triggers the most significant difference in terms of recommended item ratings. Similar approaches can be used to determine the most influencing items in other recommender types [5, 59].

Collaborative Filtering Explanations for Groups

An example of basic explanations in group-based collaborative filtering is included in POLYLENS, where the predicted rating for each group member and for the group as a whole is shown [49]. Some simple examples of how to provide explanations in the context of group-based collaborative filtering scenarios are provided in Tables 6.1 and 6.2. Both examples represent variants of the explanation approaches introduced by Herlocker et al. [30]. Table 6.1 depicts an example of an explanation that is based on the preferences (ratings) of the nearest neighbors ($NN = \bigcup\{n_{ij}\}$) of the group members u_i (for simplicity, we assume the availability of a complete set of rating data). For each recommended item t_i , the corresponding frequency distribution of the ratings of the nearest neighbors of individual group members is shown. Note that NN can represent users who are in the intersection of users who rated this item ($\{n_{11}, n_{12}, \dots\} \cap \dots \cap \{n_{m1}, n_{mk}, \dots\}$). Alternatively, NN can represent the users in the union of nearest neighbors ($\{n_{11}, n_{12}, \dots\} \cup \dots \cup \{n_{m1}, n_{mk}, \dots\}$). A related explanation can be 'users similar to members of this group rated item t as follows'.

rec. item t_i	ratings of nearest neighbors $n_{ij} \in NN$						explanation		
	u_1		u_2		u_3		bad [0 – 2]	neutral [> 2 – 3.5]	good [> 3.5 – 5]
	nn_{11}	nn_{12}	nn_{21}	nn_{22}	nn_{31}	nn_{32}			
t_1	4.2	4.9	4.3	3.5	3.2	4.8	0	2	4
t_2	3.5	2.2	2.7	3.2	2.9	3.6	0	5	1
t_3	3.8	3.1	3.7	2.8	3.4	2.6	0	4	2
t_4	4.3	4.9	4.4	4.5	4.0	4.0	0	0	6
t_5	3.7	3.9	3.2	3.5	3.6	2.9	0	3	3

Table 6.1: Collaborative filtering explanations for *aggregated predictions*, i.e., explanations based on information about the preferences (ratings) of nearest neighbors (n_{ij}) of individual group members u_i .

Table 6.2 depicts an example of an explanation that is based on the preferences of neighborhood groups gp_j of the current group gp . We assume that ratings are only available in an aggregated fashion (ratings of individual users are not available, e.g., for privacy reasons). In this context, the frequency distribution of the ratings of the nearest neighbor groups is shown for each item t_i . An explanation can contain the following text: 'groups similar to the current group rated item t as follows'.

In the given examples, explanations refer to ratings but do not take into account aggregation functions that were used (see Chapter 2). Ntoutsis et al. [47] present an approach to explain the aggregation functions in aggregated-prediction-based collaborative filtering. For example, the application of *Least Misery* (LMS) triggers explanations of type 'item y has a group score of 2.9 due to the (lowest) rating determined for user a '. A more 'group-oriented' explanation is 'item y is recommended because it avoids misery within the group'. When using *Most Pleasure*

rec. item	ratings of NN groups (gp_j)				explanation		
	gp_1	gp_2	gp_3	gp_4	bad [0 – 2]	neutral [> 2 – 3.5]	good [> 3.5 – 5]
t_1	4.2	4.9	4.3	3.5	0	1	3
t_2	1.2	2.9	3.1	1.8	2	2	0
t_3	3.5	3.8	2.9	3.3	0	3	1
t_4	4.9	4.8	4.1	4.4	0	0	4
t_5	3.7	3.3	2.4	3.9	0	2	2

Table 6.2: Collaborative filtering explanations for *aggregated models*, i.e., explanations are based on the aggregated preferences of individual group members.

(MPL), the corresponding explanation would be ‘item y has a group score of 4.8 due to the (highest) rating determined for user b ’. Finally, when using Average (AVG), explanations of type ‘item y is most similar to the ratings of users a, b , and c ’ are provided. Similar explanations can be generated for content-, constraint-, and critiquing-based recommendations. Although initial approaches have already been proposed, different ways to explain group recommendations depending on the used aggregation function(s) are an issue for future research.

Visualization of Collaborative Filtering Explanations for Groups

There are different ways to visualize a recommendation determined using collaborative filtering [30]. The frequency distributions introduced and evaluated by Herlocker et al. [30] can also be applied in the context of group recommendation scenarios. An example thereof is given in Figure 6.1, where the explanation information contained in Table 6.1 is represented graphically. Figure 6.2 depicts a similar example where an item-specific evaluation of nearest (most similar) groups is shown in terms of a frequency distribution. Alternatively, *spider diagrams* can be applied to visualize the preferences of nearest neighbors. An example is depicted in Figure 6.3. This type of representation is based on the idea of consensus-based approaches to visualize the current status of a group decision process [41, 50].

6.3 Content-based Filtering

The basis for determining recommendations in content-based filtering is the similarity between item descriptions and keywords (categories) stored in a user profile. Since the importance of keywords can differ among group members, it is important to identify those which are relevant for all group members [39]. Explanations are based on the analysis of item-related content. Examples of verbal explanations in content-based filtering are given in [5]. The authors show that keyword-style explanations can increase both the perceived trustworthiness and the transparency of rec-

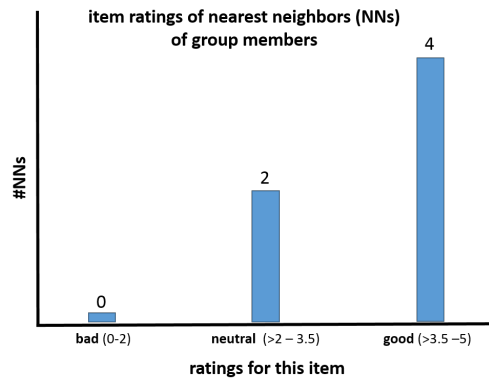


Fig. 6.1 Graphical representation of the explanation data contained in Table 6.1.




Fig. 6.2 Graphical representation of the explanation data contained in Table 6.2.




Fig. 6.3 Spider diagram for explaining aggregated models based collaborative filtering recommendations: ratings of nearest neighbor groups gp_1, \dots, gp_4 of gp for the recommended item t_4 . This representation is a variant of consensus-based interfaces discussed in [41].

ommendations. Such explanations primarily represent occurrence statistics of keywords in item descriptions (see also [14]). Gedikli et al. [28] compare different approaches to representing explanations in content-based filtering scenarios, and show that tag-cloud-based graphical representations outperform verbal approaches.

Content-based Filtering Explanations for Groups

A simple example of content-based filtering explanations for groups is depicted in Table 6.3.

category	userweights			itemweights				explanation-relevance			
	u_1	u_2	u_3	t_1	t_2	t_3	t_4	t_1	t_2	t_3	t_4
cat_1	0.05	0.1	0.15	0.1	0.1	0.2	0.3	0.01	0.01	0.02	0.03
cat_2	0.3	0.4	0.5	0.7	0.2	0.2	0.0	0.28√	0.08√	0.08	0.0
cat_3	0.15	0.25	0.2	0.1	0.4	0.2	0.3	0.02	0.08√	0.04	0.06√
cat_4	0.4	0.3	0.2	0.1	0.2	0.3	0.1	0.03	0.06	0.09√	0.03

Table 6.3: Content-based filtering explanations for *aggregated predictions*. The most explanation-relevant categories for an item t_k are marked with √.

Item categories cat_j have a user-specific weight (derived, for example, from the category weights of individual user profiles where user u_i is a member of group G). To determine the *explanation relevance* of individual categories, these weights are combined with item-individual weights (see Formula 6.1).

$$explanation-relevance(cat_j, t_k) = \frac{\sum_{u_i \in G} userweight(u_i, cat_j) \times itemweight(t_k, cat_j)}{|G|} \quad (6.1)$$

The higher the explanation-relevance of a category, the higher the category will be ranked in a list shown to the group (members). A verbal explanation related to item t_1 (Table 6.3) can be of the form ‘item t_1 is recommended since each group member is interested in category cat_2 ’. If the preference information of individual group members is not available (e.g., for privacy reasons), this explanation would be formulated as ‘item t_1 is recommended since the group as a whole is interested in category cat_2 ’. Also, more than one category can be used in such an explanation. As mentioned, category- or keyword-based explanations can also be extended with information about the most influential items [5]. This can be achieved by determining those items that trigger the most significant change in item rating predictions (if not taken into account by the recommendation algorithm).

An approach to explaining recommendations on the basis of tags is presented in Vig et al. [68]. *Tagsplanations* (explanations based on user community tags) are introduced to explain recommendations. In this context, *tag relevance* is defined as the Pearson Correlation (see Chapter 1) between item ratings and corresponding tag preference values. *Tag preference* is the relationship between the number of times a specific tag has been applied to an item compared to the total number of tags applied to the item (weighted with corresponding item ratings). In a study with MOVIELENS [44] users, the authors show that both tag relevance and tag preference help to achieve the explanation goals of *justification* (why has an item been recommended) and *effectiveness* (better decisions are made). Similar to the exam-

ple shown in Table 6.3, explanation-relevance (in this case tag relevance) is used to order a list of explanatory tags [68].

An *opinion mining* approach to generating explanations is introduced by Muhammad et al. [45]. In the context of opinion mining, features are extracted from item reviews [15] and then associated with corresponding sentiment scores. Features and corresponding sentiments are then used to generate explanations related to the *pros* and *cons* of specific items. Features are sorted into *pro* or *con* according to whether their values are above or below a predetermined threshold. If we assume, for example, a threshold of 0.4, all item features with an explanation relevance ≥ 0.4 are regarded as pros, the others are regarded as cons. Formula 6.2 represents an approach to determine the explanation-relevance of a specific feature f_i where *sentiment* represents a group preference with regard to a specific feature and *item-sentiment* represents the support of the feature by the item t_j .

$$\text{explanation-relevance}(f_i) = \text{sentiment}(f_i) \times \text{item-sentiment}(t_j, f_i) \quad (6.2)$$

Opinion mining approaches to explanations can also be extended to groups. An example of applying Formula 6.2 in the context of group recommender systems is given in Table 6.4.

group profile (gp)		item-sentiments				explanation-relevance			
feature	sentiment	t_1	t_2	t_3	t_4	t_1	t_2	t_3	t_4
f_1	0.10	0.19	0.23	0.35	0.68	0.019	0.023	0.035	0.068
f_2	0.76	0.61	0.52	0.47	0.52	0.46	0.40	0.36	0.40
f_3	0.21	0.47	0.43	0.21	0.31	0.10	0.09	0.04	0.07
f_4	0.82	0.92	0.76	0.49	0.77	0.75√	0.62√	0.40√	0.63√

Table 6.4: Opinion mining based explanations for *aggregated models*. Features f_i with the highest explanation-relevance are marked with \sqrt .

This example sketches the generation of explanations in *aggregated models* scenarios. When determining explanations in the context of *aggregated predictions*, explanation relevance could be determined for each individual user and then aggregated using an aggregation function such as *Average (AVG)* to select explanations considered most relevant for the group.

Visualization of Content-based Filtering Explanations for Groups

An alternative to list-based representations of explanations is mentioned, for example, in Gedikli et al. [28], where content-based explanations are visualized in the form of *tag-clouds*. An example of a tag-cloud-based explanation in the context of group recommendation is depicted in Figure 6.4. The used tags are related to our working example from the travel domain (see Chapter 2). In this scenario, the tag-

cloud represents an explanation based on the *aggregated preferences* of individual group members. For example, *Leo* and *Isa* like city tours. One can imagine other visual encodings in terms of shape, textures, and highlightings [35]. Tag relevance can be determined on the basis of a tag relevance estimator similar to Formula 6.1.



Fig. 6.4 Tag-cloud representation used to show the relevance of tags with regard to a specific item extended with preference information related to group members (*Isa*, *Joe*, and *Leo*).

6.4 Constraint-based Recommendation

Constraint-based recommender systems are built upon deep knowledge about items and their corresponding recommendation rules (constraints). This information serves as a basis for explaining item recommendations by analyzing reasoning steps that led to the derivation of solutions (items) [25]. Such explanations follow the tradition of AI-based expert systems [6, 26]. On the one hand, explanations are used to answer *how*-questions, i.e., questions related to the reasons behind a recommendation. A corresponding analysis is provided, for example, by Felfernig et al. [21]. *How* questions are answered in terms of showing the relationship between defined user requirements req_i and the recommended items. An example of such an explanation is '*item y is recommended, since you specified the upper price limit with 500 and you preferred light-weight cameras*' (for details see [21, 25]). Besides answering *how* questions, constraint-based recommenders help to answer *why* and *why not* questions. Explanations for the first type are used to provide insights to the user as to why certain questions have to be answered, whereas explanations for *why not* questions help a user to escape from the *no solution could be found* dilemma [23] (see also Chapter 2). Felfernig et al. [21] show that such explanations can help to increase a user's trust in the recommender application. Furthermore, explanations related to *why not* questions can increase the perception of item domain knowledge.

Explanations in Constraint-based Recommendation for Groups

Formula 6.3 represents a simple example of an approach to determine the explanation-relevance of user requirements in constraint-based recommendation scenarios for groups. A related example is depicted in Table 6.5. The assumption is that all group members have already agreed on the set of requirements $\bigcup req_j$ and each group member has also specified his/her preference in terms of an importance value. An explanation that can be provided to a group in such a context is 'requirement req_3 is considered important by the whole group'.

$$explanation-relevance(req_j) = \frac{\sum_{u_i \in G} importance(req_j, u_i)}{|G|} \quad (6.3)$$

requirement	importance			explanation relevance
	u_1	u_2	u_3	
req_1	0.2	0.3	0.4	0.3
req_2	0.5	0.4	0.1	0.33
req_3	0.3	0.3	0.5	0.37 \sqrt

Table 6.5: Explanation relevance of requirements in constraint-based recommendation (*aggregated models*). The most relevant requirement is marked with \sqrt .

The example explanation shown in Table 6.5 does not take into account causal relationships between requirements and items [25]. For example, if a group agrees that the *price* of a camera has to be below 1,000 and every camera fulfills this criteria, the price requirement does not filter out items from the itemset, so there is no causal relationship between a recommendation subset of a given itemset and the price requirement.

Combining Constraints and Utilities

Constraint-based recommendation is often *combined with* an additional mechanism that supports the ranking of candidate items (see Chapter 1). An example thereof is Multi-Attribute Utility Theory (MAUT) [69] that supports the evaluation of items in terms of a set of *interest dimensions* which can be interpreted as generic requirements. For example, in the digital camera domain, *output quality* is an interest dimension that is related to user requirements such as *resolution* and *sensor size*. Group members specify their preferences with regard to the importance of the interest dimensions dim_i . Furthermore, items t_j have different contributions with regard to these dimensions (see Table 6.6).

Similar to content-based filtering, the item-specific *explanation relevance* of individual interest dimensions can be determined on the basis of Formula 6.4 where

dimension	importance			contribution			explanation relevance		
	u_1	u_2	u_3	t_1	t_2	t_3	t_1	t_2	t_3
dim_1	0.1	0.3	0.1	0.3	0.3	0.1	0.05	0.05	0.02
dim_2	0.6	0.5	0.3	0.3	0.5	0.6	0.14	0.23 \sqrt	0.28 \sqrt
dim_3	0.3	0.2	0.6	0.4	0.2	0.2	0.15 \sqrt	0.07	0.07

Table 6.6: Explanation relevance of interest dimensions in utility-based recommendation (*aggregated predictions*). The most relevant dimension is marked with \sqrt .

imp represents the user-specific importance of an interest dimension dim_i and con the contribution of an item to dim_i .

$$explanation-relevance(dim_i, t_j) = \frac{\sum_{u_k \in G} (imp(u_k, dim_i) \times con(t_j, dim_i))}{|G|} \quad (6.4)$$

Following this approach, [7, 20, 59, 60] show how to apply utility-based approaches to the selection of evaluative arguments², i.e., arguments with the highest relevance. In this context, arguments take over the role of the previously-mentioned interest dimensions. Such an approach is provided in the INTRIGUE system [3], where recommended travel destinations are explained to groups, and arguments are chosen depending on their utility for individual group members or subgroups.

An example of an argument (as an elementary component of an explanation) for a car recommended by a constraint-based recommender is '*very energy-efficient*', where *energy-efficiency* can be regarded as an interest dimension. The contribution of an item to this interest dimension is high if, for example, the fuel consumption of a car is low. If a customer is interested in energy-efficient cars and a car is energy efficient, the corresponding argument will be included in the explanation (see the example in Table 6.6). An example explanation from another domain (financial services) is the following: '*financial service t_1 is recommended since all group members strongly prefer low-risk investments*'. Examples of interest dimensions used in this context are *risk*, *availability*, and *profit*.

Consensus in Group Decisions

Situations can occur where the preferences of individual group members become inconsistent [16, 24, 41]. In the context of group recommendation scenarios, consensus is defined in terms of disagreement between individual group members regarding item evaluations (ratings) [2].³ To provide a basis for establishing consensus, such situations have to be explained and visualized [31, 41]. In this context, diagnosis methods (see Chapters 1 and 2) can help to determine repair actions that propose changes to the current set of requirements (preferences) such that a rec-

² In line with Jameson and Smyth [33], we interpret arguments as elementary parts of explanations.

³ See also Chapter 3.

ommendation can be identified. Such repairs are able to take into account the individual preferences of group members [16]. The potential of aggregation functions (see, e.g., Table 2.2) to foster consensus in group decision making is discussed in Salamo et al. [55]. Concepts to take into account consensus in group decision making are also presented in [2, 8, 9]. In scenarios such as software requirements engineering [46], there are often misconceptions regarding the evaluation/selection of a specific requirement. For example, there could be misconceptions regarding the assignment of a requirement to a software release. An explanation in such contexts indicates possible changes of requirements (assignments) that help to restore consistency (see Chapter 2). In group-based settings, such repair-related explanations help group members understand the constraints of other group members and decide in which way their own requirements should be adapted.

User-generated Explanations

User-generated explanations are defined by a group member (typically, the creator of a decision task) to explain, for example, why a specific alternative has been selected. The impact of user-generated explanations in constraint-based group recommendation scenarios was analyzed by Stettinger et. al [58]. The creator of a decision task (prioritization decisions in the context of software requirements engineering) had to explain the decision outcome verbally. In groups where such explanations were provided, this contributed to an increased satisfaction with the final decision and an increased perceived degree of group decision support quality [58]. User-generated explanations are not limited to constraint-based recommendation. For example, crowd-sourcing based approaches are based on the similar idea of collecting explanations directly from users.

Fairness Aspects in Groups

Fair recommendations in group settings can be characterized as *recommendations without favoritism or discrimination towards specific group members*. The perceived importance of fairness, depending on the underlying item domain, has been analyzed in [17]. An outcome of this study is that in high-involvement item domains (e.g., decisions regarding new cars, financial services, and apartments), the preferred preference aggregation strategies (see Chapter 2) differ from low-involvement item domains such as restaurants and movies. The latter are often the domains of repeated group decisions (e.g., the same group selects a restaurant for a dinner every three months). Groups tend to apply strategies such as *Least Misery* (LMS), in high involvement item domains, and to prefer *Average Voting* (AVG) in low-involvement item domains. When recommending packages, the task is to recommend a set of items in such a way that individual group members perceive the recommendation as fair [56]. One interpretation of fairness stated in Serbos et al. [56] is that there are at least m items included in the package that a group member likes (see Chapter 3).

An approach to take into account fairness in *repeated group decisions* is presented by Quijano-Sanchez et al. [52], where rating predictions are adapted to achieve fairness in future recommendation settings. This adaptation also depends on the personality of a group member. For example, a group member with a strong personality who was treated less favorably last time, will be immediately compensated in the upcoming group decision (see Chapter 9). A similar interpretation of fairness is introduced in Stettinger et al. [57] where fairness is also defined in the context of repeated group decisions, i.e., decisions that repeatedly take place within the same or stable groups (groups with a low fluctuation). Fairness in this context is achieved by introducing functions that systematically adapt preference weights, i.e., group members whose preferences were disregarded recently receive higher preference weights in upcoming decisions. For example, in the context of repeated decisions (made by the same group) regarding a restaurant for a dinner, the preferences of some group members are more often taken into account than the preferences of others. In such scenarios, the preference weights of individual group members can be adapted [57] (see Formulae 6.5–6.6).

Formula 6.6 provides a *fairness* estimate per user u_i in terms of the share of the number of supported preferences in relation to the number of defined preferences. The lower the value, the less the preferences of a user (group member of group G) have been taken into account, and the lower the corresponding degree of fairness with regard to u_i . Formula 6.5 reflects an approach to increasing fairness in upcoming recommendation sessions. If the fairness (Formula 6.6) in previous sessions was lower than average, a corresponding upgrade of user-specific importance weights takes place for each dimension. For an example of adapted weights see Table 6.7.

$$imp'(u_i, dim_j) = imp(u_i, dim_j) \times \left(1 + \left(\frac{\sum_{u \in G} fair(u)}{|G|} - fair(u_i)\right)\right) \quad (6.5)$$

$$fair(u_i) = \frac{\#supported\ preferences(u_i)}{\#group\ decisions} \quad (6.6)$$


user	importance (imp)			fairness (fair)	adapted importance (imp')		
	dim_1	dim_2	dim_3		dim_1	dim_2	dim_3
u_1	0.3	0.3	0.4	4/8=0.5	0.375	0.375	0.5
u_2	0.5	0.4	0.1	6/8=0.75	0.5	0.4	0.1
u_3	0.3	0.2	0.5	8/8=1.0	0.225	0.15	0.375

Table 6.7: An example of an adaptation of individual users' weights to take *fairness* into account. In this example, the importance (*imp*) weights of user u_1 have been increased, the weights of u_2 remain the same, and the weights of user u_3 have been decreased (the preferences of u_3 have been favored in previous decisions – a visualization is given in Figure 6.6).

Visualization of Constraint-based Explanations for Groups

An example of visualizing the importance of interest dimensions with regard to a final evaluation (utility) is given in Figure 6.5. Examples of interest dimensions when evaluating, for example, financial services, are *risk*, *profit*, and *availability*.

Fig. 6.5 Visualization of the importance of interest dimensions with regard to the overall item evaluation (the importance values are based on Table 6.6 where $dim_1 = risk$, $dim_2 = profit$, and $dim_3 = availability$).



If the degree of fairness of previous group decisions has to be made transparent to the group, for example, for explaining adaptations regarding the importance weights of individual group members, this can be achieved on the basis of a visualization as depicted in Figure 6.6. An example of a related verbal explanation is the following: *'the interest dimensions favored by user u_1 have been given more consideration since u_1 was at a disadvantage in previous decisions'*.




Fig. 6.6: Visualizing the *degree of fairness* (Formula 6.6) in repeated group decisions (e.g., decisions on restaurant visits). In this example, the visualization indicates that user u_1 was at a disadvantage in previous decisions.

6.5 Critiquing-based Recommendation

To assist users in constructing and refining preferences, critiquing-based recommender systems [11] determine recommendations based on the similarity between candidate and reference items. For example, in the domain of digital cameras, related explanations focus on item attributes such as *price*, *resolution*, and *optical zoom*. *System-generated critiques* (e.g., compound critiques [42]) help to explain the relationship between the currently shown reference item and candidate items. Such explanations have been found to help educate users and increase their *trust* in the underlying recommender system [51].

Critiquing-based Explanations for Groups

User-defined critiques, i.e., critiques on the current reference item directly defined by the user, can be used for the generation of explanations for recommended items (see the example in Table 6.8).

critiques of group members				support(attribute, t_i)		
attribute	crit(u_1)	crit(u_2)	crit(u_3)	t_1	t_2	t_3
price	≤ 1.000	≤ 750	≤ 600	299(1.0)	650 (0.66)	1.200 (0.0)
res	≥ 20	≥ 18	≥ 25	24(0.66)	25 (1.0)	30 (1.0)
weight	≤ 1	≤ 2	≤ 1	1.5(0.33)	3 (0.0)	2 (0.33)
exchangeable lens	y	y	n	y(0.66)	y (0.66)	n (0.33)

Table 6.8: Critiques of group members as a basis for generating explanations for item recommendations. *Support* is defined by the share of attribute-specific critiques supported by an item t_i .

In this context, $support(attribute, t_i)$ (see Formula 6.7) indicates how often an item supports a user critique on the *attribute*. For example, item t_1 supports a critique on *price* three times since all the critiques on *price* are consistent with the price of t_1 , i.e., $support(price, t_1)=1.0$. However, $support(weight, t_1)$ is only 0.33 since the weight of t_1 is 1.5 which is inconsistent with two related critiques.

$$support(attribute, t_i) = \frac{\#supportedcritiques(attribute, t_i)}{\#critiques(attribute)} \quad (6.7)$$

On the verbal level, an explanation for item t_1 could be 'the price of camera t_1 (299) is clearly within the limits specified by the group members. As expected, it has an exchangeable lens. It has a resolution (24) that satisfies the requirements of u_1 and u_2 , however, u_3 has to accept minor drawbacks. Furthermore, the weight of the camera (1.5) is significantly higher than expected by u_1 and u_3 '.

Such explanations can be provided if the preferences of group members are known. Otherwise, explanations have to be generated on the basis of *aggregated models*, where item properties are compared with the aggregated critiques defined in the group profile.

Visualization of Critiquing-based Explanations for Groups

An example of visualizing the support of different attribute-specific critiques is given in Table 6.9. The \checkmark symbol denotes the fact that the user critique on an attribute of item t_i is supported by t_i .

user	attributes(t_1)			
	<i>price</i> = 299	<i>resolution</i> = 24	<i>weight</i> = 1.5	<i>exchangeablelens</i> = y
u_1	\checkmark	\checkmark	\times	\checkmark
u_2	\checkmark	\checkmark	\checkmark	\checkmark
u_3	\checkmark	\times	\times	\times

Table 6.9: Summarization of the support-degree of user-specific critiques on item t_1 .

6.6 Conclusions and Research Issues

In this chapter, we provided an overview of explanations that help single users and groups to better understand item recommendations. As has been pointed out in pioneering work by Jameson and Smyth [33], explanations play a crucial role in group recommendation scenarios. We discussed possibilities of explaining recommendations in the context of the basic recommendation paradigms of collaborative filtering, content-based filtering, constraint-based, and critiquing-based recommendation, taking into account specific aspects of group recommendation scenarios. In order to support a more in-depth understanding of how explanations can be determined, we provided a couple of working examples of verbal explanations and corresponding visualizations.

Although extensively analyzed in the context of single-user recommendations (see, e.g., Tintarev [61]), the generation of explanations for groups entails a couple of open research issues. Specifically, aspects of group dynamics have to be analyzed with regard to their role in generating explanations. For example, *consensus*, *fairness*, and *privacy* are major aspects – the related research question is how to define explanations that best help to achieve these goals. Some initial approaches exist to explain the application of aggregation functions in group recommendation contexts (see, e.g., Ntoutsi et al. [47]), however, a more in-depth integration of social choice theories into the generation of explanations has to be performed. This is also true on the algorithmic level, as in the context of group-based configuration (see Chap-

ter 7). In this context, the integration of information about personality and emotion into explanations has to be analyzed (see also Chapter 9). Initial related work can be found, for example, in Quijano-Sanchez et al. [53] where social factors in groups are taken into account to generate *tactful explanations*, i.e., explanations that avoid, for example, damaging friendships.

Mechanisms that help to increase the quality of group decision making processes have to be investigated [36]. For example, explanations could also be used to trigger intended behavior in group decision making such as exchange of decision-relevant information among group members [4]. Finally, explaining hybrid recommendations [37] and recommendations generated by matrix factorization (MF) approaches [1, 54] are issues for future research. Summarizing, explanations for groups is a highly relevant research area with many open issues for future work.

References

1. B. Abdollahi and O. Nasraoui. Using Explainability for Constrained Matrix Factorization. In *11th ACM Conference on Recommender Systems*, pages 79–83, Como, Italy, 2017.
2. S. Amer-Yahia, S. Roy, A. Chawla, G. Das, and C. Yu. Group Recommendation: Semantics and Efficiency. In *VLDB'09*, pages 754–765, Lyon, France, 2009.
3. L. Ardissono, A. Goy, G. Petrone, M. Segnan, and P. Torasso. INTRIGUE: Personalized Recommendation of Tourist Attractions for Desktop and Handset Devices. *Applied Artificial Intelligence: Special Issue on Artificial Intelligence for Cultural Heritage and Digital Libraries*, 17(8–9):687–714, 2003.
4. M. Atas, A. Felfernig, M. Stettinger, and 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)*, pages 368–377, Oxford, UK, 2017.
5. M. Bilgic and R. Mooney. Explaining Recommendations: Satisfaction vs. Promotion. In *ACM IUI 2005 Workshop Beyond Personalization*, pages 1–6, San Diego, CA, USA, 2005.
6. B. Buchanan and E. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, 1984.
7. G. Carenini and J. Moore. Generating and evaluating Evaluative Arguments. *Artificial Intelligence*, 170(11):925 – 952, 2006.
8. J. Castro, J. Lu, G. Zhang, Y. Dong, and L. Martínez. Opinion Dynamics-Based Group Recommender Systems. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 99:1–13, 2017.
9. J. Castro, F. Quesada, I. Palomares, and L. Martínez. A Consensus-Driven Group Recommender System. *Intelligent Systems*, 30(8):887–906, 2015.
10. S. Chang, F. Harper, L. He, and L. Terveen. CrowdLens: Experimenting with Crowd-Powered Recommendation and Explanation. In *10th International AAAI Conference on Web and Social Media (ICWSM'16)*, pages 52–61. AAAI, 2016.
11. L. Chen and P. Pu. Critiquing-based Recommenders: Survey and Emerging Trends. *User Modeling and User-Adapted Interaction (UMUAI)*, 22(1–2):125–150, 2012.
12. L. Chen and F. Wang. Explaining Recommendations Based on Feature Sentiments in Product Reviews. In *ACM IUI 2017*, pages 17–28. ACM, 2017.
13. Y. Chen. Interface and Interaction Design for Group and Social Recommender Systems. In *ACM Conference on Recommender Systems (RecSys'11)*, pages 363–366, Chicago, IL, 2011.
14. H. Cramer, V. Evers, S. Ramlal, M. Van Someren, L. Rutledge, N. Stash, L. Aroyo, and B. Wielinga. The Effects of Transparency on Trust in and Acceptance of a Content-based

- Art Recommender. *User Modeling and User-Adapted Interaction (UMUAI)*, 18(5):455–496, 2008.
15. R. Dong, M. Schaal, M. OMahony, and B. Smyth. Topic Extraction from Online Reviews for Classification and Recommendation. In *23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, pages 1310–1316. AAAI, 2013.
 16. A. Felfernig, M. Atas, T.N. Trang Tran, and M. Stettinger. Towards Group-based Configuration. In *International Workshop on Configuration 2016 (ConfWS'16)*, pages 69–72, 2016.
 17. A. Felfernig, M. Atas, T.N. Trang Tran, M. Stettinger, and 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)*, pages 335–344, Arras, France, 2017.
 18. A. Felfernig and R. Burke. Constraint-based Recommender Systems: Technologies and Research Issues. In *ACM International Conference on Electronic Commerce (ICEC08)*, pages 17–26, Innsbruck, Austria, 2008.
 19. A. Felfernig, B. Gula, G. Leitner, M. Maier, R. Melcher, S. Schippel, and E. Teppan. A Dominance Model for the Calculation of Decoy Products in Recommendation Environments. In *AISB Symposium on Persuasive Technologies*, pages 43–50, Aberdeen, Scotland, 2008.
 20. A. Felfernig, B. Gula, G. Leitner, M. Maier, R. Melcher, and E. Teppan. Persuasion in Knowledge-based Recommendation. In *3rd Intl. Conf. on Persuasive Technology*, LNCS, pages 71–82. Springer, 2008.
 21. A. Felfernig, B. Gula, and E. Teppan. Knowledge-based Recommender Technologies for Marketing and Sales. *Special issue of Personalization Techniques for Recommender Systems and Intelligent User Interfaces for the International Journal of Pattern Recognition and Artificial Intelligence (IJPRAI)*, 21(2):1–22, 2006.
 22. A. Felfernig, L. Hotz, C. Bagley, and J. Tiihonen. *Knowledge-based Configuration: From Research to Business Cases*. Elsevier/Morgan Kaufmann Publishers, 1st edition, 2014.
 23. A. Felfernig, M. Schubert, G. Friedrich, M. Mandl, M. Mairitsch, and E. Teppan. Plausible Repairs for Inconsistent Requirements. In *21st International Joint Conference on Artificial Intelligence (IJCAI'09)*, pages 791–796, Pasadena, CA, 2009.
 24. A. Felfernig, M. Schubert, and C. Zehentner. An Efficient Diagnosis Algorithm for Inconsistent Constraint Sets. *Artificial Intelligence for Engineering Design, Analysis, and Manufacturing (AIEDAM)*, 26(1):53–62, 2012.
 25. G. Friedrich. Elimination of Spurious Explanations. In *16th European Conference on Artificial Intelligence (ECAI 2004)*, pages 813–817, 2004.
 26. G. Friedrich and M. Zanker. A Taxonomy for Generating Explanations in Recommender Systems. *AI Magazine*, 32(3):90–98, 2011.
 27. E. Gansner, Y. Hu, S. Kobourov, and C. Volinsky. Putting Recommendations on the Map: Visualizing Clusters and Relations. In *ACM Conference on Recommender Systems*, pages 345–348, New York, USA, 2009.
 28. F. Gedikli, D. Jannach, and M. Ge. How should I Explain? A Comparison of Different Explanation Types for Recommender Systems. *Human Computer Studies*, 72(4):367–382, 2014.
 29. S. Gkika and G. Kekakos. The Persuasive Role of Explanations in Recommender Systems. In *2nd International Workshop on Behavior Change Support Systems (BCSS 14)*, pages 59–68, 2014.
 30. J. Herlocker, J. Konstan, and J. Riedl. Explaining Collaborative Filtering Recommendations. In *ACM Conference on Computer Supported Cooperative Work*, pages 241–250. ACM, 2000.
 31. A. Jameson. More than the Sum of its Members: Challenges for Group Recommender Systems. In *International Working Conference on Advanced Visual Interfaces*, pages 48–54, 2004.
 32. A. Jameson, S. Baldes, and T. Kleinbauer. Two Methods for Enhancing Mutual Awareness in a Group Recommender System. In *ACM Intl. Working Conference on Advanced Visual Interfaces*, pages 447–449, Gallipoli, Italy, 2004.
 33. A. Jameson and B. Smyth. Recommendation to Groups. In P. Brusilovsky, A. Kobsa, and W. Nejdl, editors, *The Adaptive Web*, volume 4321 of *Lecture Notes in Computer Science*, pages 596–627. 2007.

34. A. Jameson, M. Willemsen, A. Felfernig, M. de Gemmis, P. Lops, G. Semeraro, and L. Chen. Human Decision Making and Recommender Systems. In F. Ricci, L. Rokach, and B. Shapira, editors, *Recommender Systems Handbook, 2nd Edition*, pages 611–648. Springer, 2015.
35. E. Knutov, P. DeBra, and M. Pechenizkiy. AH 12 Years Later: A Comprehensive Survey of Adaptive Hypermedia Methods and Techniques. *New Review of Hypermedia and Multimedia*, 15(1):5–38, 2009.
36. J. Konstan and J. Riedl. Recommender Systems: From Algorithms to User Experience. *User Modeling and User-Adapted Interaction (UMUAI)*, 22(1):101–123, 2012.
37. P. Kouki, J. Schaffer, J. Pujara, J. O’Donovan, and L. Getoor. User Preferences for Hybrid Explanations. In *11th ACM Conference on Recommender Systems*, pages 84–88, Como, Italy, 2017.
38. B. Lamche, U. Adigüzel, and W. Wörndl. Interactive Explanations in Mobile Shopping Recommender Systems. In *8th ACM Conference on Recommender Systems, Joint Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS’14)*, pages 14–21, Foster City, Silicon Valley, California, USA, 2014.
39. H. Lieberman, N. Dyke, and A. Vivacqua. Let’s Browse: A Collaborative Web Browsing Agent. In *4th International Conference on Intelligent User Interfaces*, pages 65–68, Los Angeles, CA, USA, 1999.
40. W. Lin, S. Alvarez, and C. Ruiz. Efficient Adaptive-Support Association Rule Mining for Recommender Systems. *Data Mining and Knowledge Discovery*, 6:83–105, 2002.
41. N. Mahyar, W. Liu, S. Xiao, J. Browne, M. Yang, and S. Dow. Consensus: Visualizing points of disagreement for multi-criteria collaborative decision making. In *ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 17–20. ACM, 2017.
42. K. McCarthy, J. Reilly, L. McGinty, and B. Smyth. On the Dynamic Generation of Compound Critiques in Conversational Recommender Systems. In *International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, pages 176–184. Springer, 2004.
43. K. McCarthy, J. Reilly, L. McGinty, and B. Smyth. Thinking Positively - Explanatory Feedback for Conversational Recommender Systems. In *European Conference on Case-Based Reasoning (ECCBR-04) Explanation Workshop*, pages 1–10, 2004.
44. B. Miller, I. Albert, S. Lam, J. Konstan, and J. Riedl. MovieLens Unplugged: Experiences with a Recommender System on Four Mobile Devices. In E. O’Neill, P. Palanque, and P. Johnson, editors, *People and Computers XVII Designing for Society*, pages 263–279. Springer, London, 2004.
45. K. Muhammad, A. Lawlor, and B. Smyth. A Live-User Study of Opinionated Explanations for Recommender Systems. In *21st International Conference on Intelligent User Interfaces (IUI 2016)*, pages 256–260. ACM, 2016.
46. G. Ninaus, A. Felfernig, M. Stettinger, S. Reiterer, G. Leitner, L. Weninger, and W. Schanil. INTELLIREQ: Intelligent Techniques for Software Requirements Engineering. In *Prestigious Applications of Intelligent Systems Conference (PAIS)*, pages 1161–1166, 2014.
47. E. Ntoutsis, K. Stefanidis, K. Norvag, and H. Kriegel. Fast Group Recommendations by Applying User Clustering. In *ER 2012*, volume 7532 of *LNCS*, pages 126–140, 2012.
48. I. Nunes and D. Jannach. A Systematic Review and Taxonomy of Explanations in Decision Support and Recommender Systems. *User Modeling and User-Adapted Interaction (UMUAI)*, 2017.
49. M. O’Connor, D. Cosley, J. Konstan, and J. Riedl. PolyLens: A Recommender System for Groups of Users. In *7th European Conference on Computer Supported Cooperative Work*, pages 199–218, 2001.
50. I. Palomares, L. Martinez, and F. Herrera. MENTOR: A Graphical Monitoring Tool of Preferences Evolution in Large-Scale Group Decision Making. *Knowledge-Based Systems*, 58:66–74, 2014.
51. P. Pu and L. Chen. Trust-inspiring Explanation Interfaces for Recommender Systems. *Knowledge-based Systems*, 20(6):542–556, 2007.
52. L. Quijano-Sanchez, J. Recio-García, B. Díaz-Agudo, and G. Jiménez-Díaz. Social Factors in Group Recommender Systems. *ACM Trans. on Intell. Sys. and Tech.*, 4(1):8:1–8:30, 2006.

53. L. Quijano-Sanchez, C. Sauer, J. Recio-García, and B. Díaz-Agudo. Make it Personal: A Social Explanation System applied to Group Recommendations. *Expert Systems with Applications*, 76:36–48, 2017.
54. B. Rastegarpanah, M. Crovella, and K. Gummadi. Exploring Explanations for Matrix Factorization Recommender Systems, FatRec Workshop. In *11th ACM Conference on Recommender Systems*, Como, Italy, 2017.
55. M. Salamo, K. McCarthy, and B. Smyth. Generating Recommendations for Consensus Negotiation in Group Personalization Services. *Personal and Ubiquitous Computing*, 16(5):597–610, 2012.
56. D. Serbos, S. Qi, N. Mamoulis, E. Pitoura, and P. Tsaparas. Fairness in Package-to-Group Recommendations. In *WWW'17*, pages 371–379. ACM, 2017.
57. M. Stettinger. CHOICLA: Towards Domain-independent Decision Support for Groups of Users. In *8th ACM Conference on Recommender Systems*, pages 425–428, Foster City, Silicon Valley, California, USA, 2014.
58. M. Stettinger, A. Felfernig, G. Leitner, and S. Reiterer. Counteracting Anchoring Effects in Group Decision Making. In *23rd Conference on User Modeling, Adaptation, and Personalization (UMAP'15)*, volume 9146 of *LNCIS*, pages 118–130, Dublin, Ireland, 2015.
59. P. Symeonidis, A. Nanopoulos, and Y. Manolopoulos. Providing Justifications in Recommender Systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 38:1262–1272, 2008.
60. J. Teze, S. Gottifredi, A. Garcia, and G. Simari. Improving Argumentation-based Recommender Systems through Context-adaptable Selection Criteria. *Journal of Economic Perspectives*, 42(21):8243–8258, 2015.
61. N. Tintarev. *Explaining Recommendations*. University of Aberdeen, 2009.
62. N. Tintarev and J. Masthoff. Designing and Evaluating Explanations for Recommender Systems. In *Recommender Systems Handbook*, pages 479–510. 2011.
63. N. Tintarev and J. Masthoff. Evaluating the Effectiveness of Explanations for Recommender Systems. *User Modeling and User-Adapted Interaction*, 22(4–5):399–439, 2012.
64. N. Tintarev and J. Masthoff. Explaining Recommendations: Design and Evaluation. In F. Ricci, L. Rokach, and B. Shapira, editors, *Recommender Systems Handbook, 2nd Edition*, pages 353–382. Springer, 2015.
65. N. Tintarev, J. O'Donovan, and A. Felfernig. Human Interaction with Artificial Advice Givers. *ACM Transactions on Interactive Intelligent Systems*, 6(4):1–10, 2016.
66. T. Ulz, M. Schwarz, A. Felfernig, S. Haas, A. Shehadeh, S. Reiterer, and M. Stettinger. Human Computation for Constraint-based Recommenders. *Journal of Intelligent Information Systems (JIIS)*, 49(1):37–57, 2017.
67. K. Verbert, D. Parra, P. Brusilovsky, and E. Duval. Visualizing Recommendations to Support Exploration, Transparency and Controllability. In *International Conference on Intelligent User Interfaces (IUI'13)*, pages 351–362, New York, NY, USA, 2013.
68. J. Vig, S. Sen, and J. Riedl. Tagsplanations: Explaining Recommendations Using Tags. In *ACM IUI 2009*, pages 47–56, Sanibel Island, FL, USA, 2009. ACM.
69. D. Winterfeldt and W. Edwards. *Decision Analysis and Behavioral Research*. Cambridge University Press, 1986.