

Chapter 1

Decision Tasks and Basic Algorithms



Alexander Felfernig, Müslim Atas, and Martin Stettinger

Abstract Recommender systems are *decision support systems* helping users to identify one or more items (solutions) that fit their wishes and needs. The most frequent application of recommender systems nowadays is to propose items to *individual users*. However, there are many scenarios where a *group of users* should receive a recommendation. For example, think of a group decision regarding the *next holiday destination* or a group decision regarding a *restaurant to visit for a joint dinner*. The goal of this book is to provide an introduction to *group recommender systems*, i.e., recommender systems that determine recommendations for groups. In this chapter, we provide an introduction to basic types of recommendation algorithms for individual users and characterize related decision tasks. This introduction serves as a basis for the introduction of group recommendation algorithms in Chap. 2.

1.1 Introduction

A recommender system is *a specific type of advice-giving or decision support system that guides users in a personalized way to interesting or useful objects in a large space of possible options or that produces such objects as output* [14, 17, 23, 29, 43, 66, 76, 81, 82]. A decision problem/task emerges if a person or a group of persons have an idea about a desired state [33]. If there are different options to achieve the desired state, the *decision task* to be solved is to identify items or actions that help to approach the target state in a suitable fashion. Recommender systems can provide help in such a context by trying to find the suitable items or actions that help to best reach the envisioned target. Arriving at a choice can be seen as the result of a collaboration between the user and the recommender system. Recommender systems support “good” choices within reasonable time spans including corresponding justifications provided in terms of explanations [42, 80].

There are different ways in which a recommender can support users in decision making processes. First, it can act as a *supporter* to figure out *candidate items*, i.e., a large number of alternatives is reduced to a so-called *consideration set*—selecting the favorite option is left to the user. Second, the recommender can help the user select from the items in the consideration set, for example, by representing them in convenient ways and providing explanations of why they have been recommended. Only in “extreme” cases is the *decision making authority taken over* by the recommender itself. Examples include music recommendation in a fitness studio, the recommendation of information units on a public display, and the automated adaptation of parameter settings such as light intensity in a smarthome [50].

An example of a single-user recommendation scenario is the following: when navigating an online sales platform with the goal to find a book related to the topic of, for example, *deep learning*, a recommender system will identify related books and propose these to the user of the online platform. In this scenario, the envisioned target state is to find a suitable book on the mentioned topic and the options to achieve this state are the different existing books on the topic of *deep learning*. Finding a book in an online sales platform is typically a *single user decision scenario*. However, there are also scenarios where a group of users has to make a decision. In this context, a recommender system must take into account the potentially conflicting preferences of different group members. Such a situation makes the recommendation task different and often more challenging.

The main focus of this book is *recommendation techniques that provide help in scenarios where a group of users is engaged*. In many scenarios, the presentation of recommendations to groups is a more natural approach than trying to address individual users [40, 56, 58]. For example, music recommendations in fitness studios have to take into account the preferences of all individuals currently present in the studio [59]. Stakeholders in a software project have to establish agreement regarding the requirements/features that have to be developed within the scope of the next release [64]. Personnel decisions are often taken in groups, i.e., a group has to decide which job applicant will be hired [78]. Groups of friends have to decide about the hotel for the next summer holidays or a skiing resort for the next winter holidays [39, 61]. A public display should be personalized in order to be able to display information to persons currently in the surrounding [40]. Finally, travel groups should receive a personalized museum guidance in such a way that the personal preferences of group members are fulfilled [28, 47].

1.2 Characteristics of Decision Tasks

Decision tasks differ with regard to various aspects [69]. In the following, we introduce *basic dimensions of decision tasks* (see [33]) which will help to better understand which decisions are supported by group recommenders (see Table 1.1).¹

¹As mentioned in [33], this characterization of decision tasks is not complete but a good basis for discussing properties relevant for group recommenders.

Table 1.1 Characteristics of decision tasks [33]

Dimension	Characteristic
Complexity	low .. high
Structuredness	low .. high
Decision type	choice .. design
Sentiment	opportunity .. threat
Dependence	yes .. no
Level	original .. meta
Actor	person .. group

Complexity We interpret complexity of decision tasks in terms of the *number of decision dimensions* and the *degree of item involvement* [33, 38]. Depending on the underlying item domain, humans will invest more or less time to come to a decision, i.e., to achieve an acceptable trade-off between decision effort and accuracy [67]. Items with higher related decision efforts are often denoted as *high-involvement items* whereas items with less related decision efforts are denoted as *low-involvement items* [70]. Suboptimal decisions have a much higher negative impact in the context of high-involvement items. For example, when purchasing an *apartment*, a suboptimal decision manifests in search efforts for a new apartment, unnecessary payments, relocation costs, and additional time efforts [27]. In contrast, risks related to the purchasing of a low-quality book are negligible—in the worst case, the user will provide negative feedback on the book and try to find other alternatives that better fit his/her wishes and needs. When purchasing a book, the number of decision dimensions is low—examples of dimensions are *price* and *quality*. The number of decision dimensions of *apartments* is much higher (e.g., price, quality of public transport, neighborhood, schools in the neighborhood, etc.). Aspects that further increase decision complexity especially in group decision scenarios are, for example, *contradicting preferences of group members*, *personal relationships*, *personality factors*, and *emotion-related aspects* (see Chap. 9).

Structuredness We interpret structuredness of decision tasks as the degree to which underlying processes and decision policies are defined. Decision tasks are often characterized by undefined processes and related decision policies are not pre-defined but developed and adapted in the course of the decision process. If a group of users has to decide on a holiday destination for the next summer, a recommender system can propose alternative destinations but it is unclear which of the alternatives will be chosen by the group. The final decision is something that has to be made by group members (or an individual user) and is in many cases not handled by the underlying decision support environment. There are exceptions to the rule, for example, music recommendations in fitness centers and information units shown on public displays. Specific decision types follow a *formalized process*. For example, electoral systems are defined by precise rules that determine how elections and referendums have to be conducted (the process) and how the results are determined (decision making).

Decision Type Decision tasks are often defined on the basis of known alternatives or parameters out of which one or more alternatives (values) should be selected. If alternatives (parameters) are predefined, the underlying decision task can be regarded as a basic *choice task* [42]. Choice problems are considered as central application area for recommender systems [42]. The other extreme are so-called *design tasks*, where alternatives are not predefined but created throughout a decision process. Design tasks are often related to creative acts where persons develop ideas and solutions. In “pure” design scenarios, the application of recommendation technologies is not widespread [75]. However, there are many scenarios located in-between basic choice tasks and “pure” design tasks. For example, knowledge-based configuration [25] can be considered as a simpler type of design task where a solution is identified (configured) out of a set of pre-defined component types. In this context, the alternatives (parameter values) are known beforehand; due to the large option space, not all potential alternatives can be enumerated for performance reasons—billions of alternatives would have to be managed and the corresponding recommendation algorithms would become inefficient [13] (see Chap. 7).

Sentiment Decision support with included recommendation support is very often *opportunity-related*, i.e., the goal is related to an opportunity and the best solution to achieve a goal should be identified. Examples thereof are purchasing a book to better understand a certain topic or choosing a holiday destination to spend unforgettable days somewhere abroad. A similar argument holds for item domains such as songs, digital equipment, food, and financial services. Certainly, decision problems also exist in contexts where alternative outcomes can be considered as negative ones. For example, choosing between alternative options to liquidate a company—in this scenario, every outcome can be considered as a negative one (the company gets liquidated). However, recommender systems can help to minimize damage, for example, on the basis of a structured utility analysis [79].

Dependence We regard decision tasks as dependent if the outcome of a decision has an impact on another decision. For example, the purchase of a movie typically does not require a follow-up decision regarding the purchase of the next movie or different item. Dependent decision tasks occur when one decision at an earlier point of time leads to follow-up decision tasks at a later point of time. An example of such a decision task is *requirements release planning* where for each software requirement it has to be decided when the requirement should be implemented [64]. Consequently, decisions taken in early phases of a software project can have an impact on or trigger decisions later in the project. For example, a decision that a requirement should be implemented could trigger decisions regarding additional resources in order to be able to provide the promised software functionalities in time.

Level We can differentiate between *original decisions* operating on the object level and decisions on the *meta-level* [33]. The first type is omni-present in many recommendation-supported scenarios—the underlying task is to identify and choose items of relevance. In contrast, meta-decisions are decisions about the qualities of

a decision process and the way decisions are taken. For example, a group could decide to use *majority voting* when it comes to the election of the next chairman. A meta-decision in this context is to decide about the election formalism—related alternatives can be, for example, *relative majority and a single-shot election* or *absolute majority in a potentially multi-level election process* [51]. In many decision scenarios—especially in the context of group decision making—recommendations have an advisory function but are not considered imperative.

Actor Many recommendation approaches support individual decision making where recommendation algorithms are focusing on determining recommendations for individual users. The focus of this book are recommender systems that support *decision making for groups of users*. The following types of groups are introduced in [7]: (1) *established groups* with shared and long-term common interests (e.g., *conference committees* taking decisions about conference venues or *families* taking a decision about purchasing a house), (2) *occasional groups* with a common aim in a particular moment (e.g., a group of persons jointly participating in a museum tour), (3) *random groups* (e.g., persons in a fitness center or persons around a public display), and (4) *automatically identified groups* where individuals with similar preferences have to be grouped (e.g., distribution of seminar papers to students and distribution of conference papers to reviewers).

1.3 Recommendation Algorithms for Individual Users

Recommender systems [43, 58] propose items of potential interest to an individual user or a group of users.² They are applied in item domains such as books [52], web sites [68], financial services [16], and software artifacts [20, 24]. In the following, we introduce *collaborative filtering* [31, 48], *content-based filtering* [68], *constraint-based* [8, 14], *critiquing-based* [10, 11], and *hybrid recommendation* [9] which are basic recommendation approaches. The items in Table 1.2 (travel destinations) serve as examples to demonstrate how basic recommendation algorithms operate. In Chap. 2, we show how these approaches can be integrated into corresponding group recommendation scenarios.

Collaborative Filtering

Collaborative filtering is based on the idea of word-of-mouth promotion where opinions of relatives and friends play a major role when taking a decision [12, 48, 52]. In online scenarios, family members and friends are replaced by *nearest neighbors* who

²Parts of this section are based on a discussion of recommendation technologies given in [24].

Table 1.2 Example set of travel destinations

Travel destination (item)	Name	Beach	City tours	Nature	Entertainment
t_1	Vienna		x		x
t_2	Yellowstone			x	
t_3	New York	x	x		x
t_4	Blue Mountains			x	
t_5	London		x		x
t_6	Beijing		x		x
t_7	Cape Town	x	x	x	x
t_8	Yosemite			x	
t_9	Paris		x		x
t_{10}	Pittsburgh		x		x

These items will be used across the following sections for demonstration purposes. Each travel destination is described by a set of meta-characteristics (categories), for example, travel destination *Yellowstone* is famous for its experience of *nature*

are users with preferences similar to the ones of the current user. In collaborative filtering, a *user* \times *item* rating specifies to which extent a user likes an item. *Rating predictions* are determined by a recommender algorithm to estimate the extent a user will like an item he/she did not consume/evaluate up to now. A collaborative filtering recommender first determines k nearest neighbors ($k - NN$).³ The preferences of nearest neighbors are used to extrapolate future item ratings of the current user. A *user* \times *item* rating matrix that will be used in the following for explaining collaborative filtering is shown in Table 1.3. In this example, all users u_i visited travel destinations and provided a corresponding rating. In collaborative filtering, *user* \times *item* ratings serve as input for the recommender.

Collaborative filtering identifies the k -nearest neighbors of the current user u_a (see Formula 1.1)⁴ and—based on the nearest neighbors—calculates a prediction for the current user’s rating. When applying Formula 1.1, user u_2 is identified as the nearest neighbor of user u_a (see also Table 1.3). The similarity between u_a and another user u_x can be determined, for example, using the Pearson correlation coefficient [43] (see Formula 1.1) where TD_c is the set of items that have been rated by both users (u_a and u_x), r_{x,t_i} is the rating of user x for item t_i , and \bar{r}_x is the average rating of user x . Similarity values resulting from Formula 1.1 can take values on a scale of $[-1.. + 1]$. Sometimes, “neighbor” users with low or negative correlations with the current user are filtered out [1].

³We focus on *user-based collaborative filtering* which is a memory-based approach that—in contrast to model-based ones—operates on an uncompressed version of a user/item matrix [6, 12].

⁴We assume $k = 2$ in our example.

Table 1.3 Example of collaborative filtering rating matrix: travel destinations (items) t_i and ratings (we assume a rating scale of 1..5)

Item	Name	u_1	u_2	u_3	u_4	u_a
t_1	Vienna	5.0			4.0	
t_2	Yellowstone	4.0				
t_3	New York		3.0	4.0	3.0	
t_4	Blue Mountains		5.0	5.0		4.0
t_5	London			3.0		
t_6	Beijing		4.5	4.0		4.0
t_7	Cape Town	4.0				
t_8	Yosemite		2.0			
t_9	Paris				3.0	
t_{10}	Pittsburgh				5.0	3.0
	Average rating (\bar{r}_x)	4.33	3.625	4.0	3.75	3.67

Table 1.4 Similarity between user u_a and the users $u_j \neq u_a$ determined based on Formula 1.1

	u_1	u_2	u_3	u_4
u_a	–	0.97	0.70	–

If the number of commonly rated items is below 2, no similarity between the two users is calculated

$$\text{similarity}(u_a, u_x) = \frac{\sum_{t_i \in TD_c} (r_{a,t_i} - \bar{r}_a) \times (r_{x,t_i} - \bar{r}_x)}{\sqrt{\sum_{t_i \in TD_c} (r_{a,t_i} - \bar{r}_a)^2} \times \sqrt{\sum_{t_i \in TD_c} (r_{x,t_i} - \bar{r}_x)^2}} \quad (1.1)$$

The similarity values for u_a calculated based on Formula 1.1 are shown in Table 1.4. For the purpose of our example, we assume the existence of at least two items per user pair (u_i, u_j) ($i \neq j$) in order to be able to determine a similarity. This criterion holds for users u_2 and u_3 .

A challenge when determining the similarity between users is the *sparsity* of the rating matrix. Users typically provide ratings for only a very small subset of the offered items. For example, given a large movie dataset that contains thousands of entries, a user will typically be able to rate only a few dozen. One approach to this problem is to take into account the number of commonly rated items as a *correlation significance* [37], i.e., the higher the number of commonly rated items, the higher is the significance of the corresponding correlation. For further information regarding the handling of sparsity, we refer to [37, 43].

The information about the set of users with a rating behavior similar to that of the current user (nearest neighbors NN) is the basis for predicting the rating of user u_a for an item t that has so far not been rated by u_a (see Formula 1.2).

$$\text{prediction}(u_a, t) = \hat{r}(u_a, t) = \bar{r}_a + \frac{\sum_{u_j \in NN} \text{similarity}(u_a, u_j) \times (r_{j,t} - \bar{r}_j)}{\sum_{u_j \in NN} \text{similarity}(u_a, u_j)} \quad (1.2)$$

Table 1.5 Collaborative filtering based recommendations (predictions) for items that have not been rated by user u_a up to now

	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}
u_2	–	–	3.0	5.0	–	4.5	–	2.0	–	–
u_3	–	–	4.0	5.0	3.0	4.0	–	–	–	–
u_a	–	–	–	4.0	–	4.0	–	–	–	3.0
Prediction for u_a	–	–	3.30 ✓	–	2.66	–	–	2.04	–	–

Based on the ratings of the nearest neighbors of u_a , we are able to determine a prediction for u_a (see Table 1.5). The nearest neighbors of u_a are assumed to be u_2 and u_3 (see Table 1.4). The travel destinations rated by the nearest neighbors but not rated by u_a are t_3 , t_5 , and t_8 . Due to the determined predictions (Formula 1.2), item t_3 would be ranked higher than the items t_5 and t_8 in a recommendation list. For a discussion of advanced collaborative recommendation approaches, we refer the reader to [49, 74].

Content-Based Filtering

This approach is based on the assumption of monotonic personal interests [68]. For example, users interested in *political news* are typically not changing their interest profile from one day to another. On the contrary, they will also be interested in the topic in the (near) future. In online scenarios, content-based recommenders are applied, for example, when it comes to the recommendation of websites [68].

Content-based filtering is based on (a) a set of users and (b) a set of categories (or keywords) that have been assigned to (or extracted from) the set of items. It compares the content of already consumed items with new items, i.e., it finds items that are similar to those already consumed (positively rated) by the user (u_a). The basis for determining such a similarity are *keywords* extracted from the item content descriptions (e.g., keywords extracted from news articles) or *categories* if items have been annotated with the relevant meta-information. Readers interested in the principles of keyword extraction are referred to [43]. In this book, we focus on content-based recommendation which exploits item categories (see Table 1.2).

Content-based filtering will now be explained using Tables 1.2, 1.6, and 1.7. Table 1.2 provides an overview of the relevant items and the assignments of items to categories. Table 1.6 provides information on which categories are of relevance for our example users. For example, user u_1 is interested in items related to all categories. Since user u_a rated the items t_4 , t_6 , and t_{10} (see Table 1.3), we can infer that u_a is interested in the categories *City Tours*, *Nature*, and *Entertainment* (see Table 1.6) where items related to the categories *City Tours* and *Entertainment* have been evaluated twice and items related to *Nature* have been evaluated once by u_a .

If we are interested in an item recommendation for the user u_a , we have to search for those items which are most similar to the items that have already been consumed

Table 1.6 Example category interests of users: users $u_1 - u_3$ visited travel destinations that cover all available categories

Category	u_1	u_2	u_3	u_4	u_a
Beach	x	x	x	x	—
City tours	x	x	x	x	x
Nature	x	x	x	—	x
Entertainment	x	x	x	x	x

If a user accessed an item at least once with a rating ≥ 3 (see Table 1.3), it is inferred that the user is interested in this item

Table 1.7 Example of content-based filtering: user u_a has already consumed the items $t_4, t_6,$ and t_{10} (see Table 1.3)

Item	Rating (u_a)	Name	Beach	City tours	Nature	Entertainment	Similarity(u_a, t_i)
t_1		Vienna	—	x	—	x	$\frac{4}{5}$
t_2		Yellowstone	—	—	x	—	$\frac{1}{2}$
t_3		New York	x	x	—	x	$\frac{2}{3}$
t_4	4.0	Blue Mountains	—	—	x	—	—
t_5		London	—	x	—	x	$\frac{4}{5}$
t_6	4.0	Beijing	—	x	—	x	—
t_7		Cape Town	x	x	x	x	$\frac{6}{7}$ ✓
t_8		Yosemite	—	—	x	—	$\frac{1}{2}$
t_9		Paris	—	x	—	x	$\frac{4}{5}$
t_{10}	3.0	Pittsburgh	—	x	—	x	—
user u_a				x	x	x	

The item most similar (see Formula 1.3) to the preferences of u_a is t_7

(evaluated) by u_a . This relies on the simple similarity metric shown in Formula 1.3 (Dice coefficient which is a variation of the Jaccard coefficient “intensively” taking into account category commonalities—see also [43]). The major difference to the similarity metrics introduced in the context of collaborative filtering is that similarity is measured using categories (in contrast to ratings).

$$similarity(u_a, item) = \frac{2 * |categories(u_a) \cap categories(item)|}{|categories(u_a)| + |categories(item)|} \tag{1.3}$$

Constraint-Based Recommendation

Compared to the approaches of collaborative filtering and content-based filtering, *constraint-based recommendation* [14, 16, 63] does not primarily rely on item ratings and textual item descriptions but on deep knowledge about the offered items. Such deep knowledge (semantic knowledge [16]) describes an item in more detail and thus allows for a different recommendation approach (see Table 1.8).

Table 1.8 Slightly adapted travel destinations described based on *season* (digit 1 indicates a recommended season and 0 indicates a non-recommended one; seasons start with *spring*), associated *topics*, and average user rating (*eval*)

Item	Name	Season	Topics	Eval
t_1	Vienna	1110	City tours, entertainment	4.5
t_2	Yellowstone	1110	Nature	4.0
t_3	New York	1011	City tours, entertainment	3.3
t_4	Blue Mountains	1001	Nature	5.0
t_5	London	1010	City tours, entertainment	3.0
t_6	Beijing	1010	City tours, entertainment	4.7
t_7	Cape Town	1111	Beach, city tours, nature, entertainment	4.0
t_8	Yosemite	1110	Nature	2.0
t_9	Paris	1011	City tours, entertainment	3.0
t_{10}	Pittsburgh	1010	City tours	5.0

Constraint-based recommendation relies on (a) a set of rules (constraints) and (b) a set of items. Depending on the user requirements (a set of search criteria), rules (constraints) describe which items should be recommended. The current user (u_a) articulates his/her requirements (preferences) in terms of item property specifications which are internally represented as rules (constraints). In our example, constraints are represented solely by user requirements, no further constraint types are included. An example of a constraint is the following: *topics = city tours* (the user is primarily interested in travel destinations allowing city tours). For a detailed discussion of further constraint types, we refer the reader to [16]. Constraints are interpreted and the resulting items are presented to the user. A detailed discussion of reasoning mechanisms that are used in constraint-based recommendation can be found in [14, 19, 22]. In order to determine a recommendation in a constraint-based recommendation scenario, a *recommendation task* has to be solved.

Definition (Recommendation Task) A recommendation task can be defined by the tuple (R, I) where R represents a set of user requirements and I represents a set of items (in our case: travel destinations $t_i \in I$). The goal is to identify those items in I which fulfill the given user requirements (preferences).

A solution for a recommendation task (also denoted as recommendation) can be defined as follows.

Definition (Solution for a Recommendation Task) A solution for a recommendation task (R, I) is a set $S \subseteq I$ such that $\forall t_i \in S : t_i \in \sigma_{(R)} I$ where σ is the selection operator of a conjunctive query [19], R represents a set of selection criteria (represented as constraints), and I represents an item table (see, e.g., Table 1.8). If we want to restrict the set of item properties shown to the user in a result set (recommendation), we have to additionally include projection criteria π as follows: $\pi_{(attributes(I))}(\sigma_{(R)} I)$.

In our example, we show how to determine a solution for a given recommendation task based on a conjunctive query where user requirements are used

Table 1.9 Travel destinations described with regard to the dimensions *security* (high evaluation represents a high security), *attractiveness* (high evaluation represents a high attractiveness), and *crowdedness* (high evaluation represents a low crowdedness)

Item	Name	Security	Attractiveness	Crowdedness
t_1	Vienna	5	5	2
t_2	Yellowstone	4	4	4
t_3	New York	3	5	1
t_4	Blue Mountains	4	3	5
t_5	London	3	4	1
t_6	Beijing	3	3	1
t_7	Cape Town	2	3	3
t_8	Yosemite	4	4	4
t_9	Paris	3	5	1
t_{10}	Pittsburgh	3	3	3

For example, *security* = 5 for the item *Vienna* indicates the highest possible *contribution* to the dimension *security* (scale 1..5)

as selection criteria (constraints) on an item table I . If we assume that the user requirements are represented by the set $R = \{r_1 : season = winter, r_2 : topics = city\ tours\}$ and the item table I consists of the elements shown in Table 1.8, then $\pi_{(item)}(\sigma_{(season=winter \wedge topics=city\ tours)}I) = \{t_3, t_7, t_9\}$, i.e., these three items are consistent with the given set of requirements.

Ranking Items Up to now we know which items can be recommended to a user. One widespread approach to rank items is to define a utility scheme which serves as a basis for the application of *Multi Attribute Utility Theory* (MAUT).⁵ Items can be evaluated and ranked with respect to a set of *interest dimensions*. In travel destinations, example interest dimensions are *security* (security level of the travel destination), *attractiveness* (estimated attractiveness of the travel destination), and *crowdedness* (degree of crowdedness of the travel destination). The first step to establish a MAUT scheme is to relate the interest dimensions with the given set of items. An example thereof is shown in Table 1.9 where the city of *Vienna* receives the highest evaluations with regard to the dimensions *security* and *attractiveness* and a low evaluation with regard to *crowdedness*.

We are now able to determine the user-specific utility of each individual item. The calculation of *item* utilities for a specific user u_a can be based on Formula 1.4.

$$utility(u_a, item) = \sum_{d \in Dimensions} contribution(item, d) \times weight(u_a, d) \quad (1.4)$$

If we assume that the current user u_a assigns a weight of 0.6 to the dimension *security* ($weight(u_a, security)=0.6$), a weight of 0.3 to the dimension *attractiveness*

⁵A detailed discussion of MAUT in constraint-based recommendation is given in [2, 16, 18].

Table 1.10 Item-specific utility for user u_a ($utility(u_a, t_i)$) assuming the personal preferences $\{\text{weight}(u_a, \text{security})=0.6, \text{weight}(u_a, \text{attractiveness})=0.3, \text{weight}(u_a, \text{crowdedness})=0.1\}$, item t_1 has the highest utility for user u_a

Item	Security	Attractiveness	Crowdedness	Utility
t_1	3.0	1.5	0.2	4.7 ✓
t_2	2.4	1.2	0.4	4.0
t_3	1.8	1.5	0.1	3.4
t_4	2.4	0.9	0.5	3.8
t_5	1.8	1.2	0.1	3.1
t_6	1.8	0.9	0.1	2.8
t_7	1.2	0.9	0.3	2.4
t_8	2.4	1.2	0.4	4.0
t_9	1.8	1.5	0.1	3.4
t_{10}	1.8	0.9	0.3	3.0

($\text{weight}(u_a, \text{attractiveness})=0.3$), and a weight of 0.1 to the dimension *crowdedness* ($\text{weight}(u_a, \text{crowdedness})=0.1$), then the user-specific utilities of the individual items (t_i) are the ones shown in Table 1.10.

Managing Inconsistencies In constraint-based recommendation scenarios, we have to deal with situations where no solution (recommendation) can be identified for a given set of user requirements, i.e., $\sigma_{(R)}I = \emptyset$. In such situations, we are interested in proposals for requirements changes such that a solution can be found. For example, if a user is interested in travel destinations for *entertainment* with an overall *evaluation of 5.0* in the summer season, then no solution can be provided for the given set of requirements $R = \{r_1 : \text{season} = \text{summer}, r_2 : \text{topics} = \text{entertainment}, r_3 : \text{eval} = 5.0\}$.

User support in such situations can be based on the concepts of conflict detection [44] and model-based diagnosis [13, 15, 72]. A conflict (or conflict set) with regard to an item set I in a given set of requirements R can be defined as follows.

Definition (Conflict Set) A conflict set is a set $CS \subseteq R$ such that $\sigma_{(CS)}I = \emptyset$. CS is minimal if there does not exist a conflict set CS' with $CS' \subset CS$.

In our example, we are able to determine the following minimal conflict sets CS_i : $CS_1 : \{r_1, r_3\}$, $CS_2 : \{r_2, r_3\}$. We will not discuss algorithms that support the determination of minimal conflict sets but refer the reader to the work of Junker [44] who introduces a divide-and-conquer based algorithm with a logarithmic complexity in terms of the needed number of consistency checks.

Based on the identified minimal conflict sets, we are able to determine the corresponding (minimal) diagnoses. A diagnosis for a given set of requirements which is inconsistent with the underlying item table can be defined as follows.

Definition (Diagnosis) A diagnosis for a set of requirements $R = \{r_1, r_2, \dots, r_n\}$ is a set $\Delta \subseteq R$ such that $\sigma_{(R-\Delta)}I \neq \emptyset$. A diagnosis Δ is minimal if there does not exist a diagnosis Δ' with $\Delta' \subset \Delta$.

In other words, a diagnosis (hitting set) is a minimal set of requirements that have to be deleted from R such that a solution can be found for $R - \Delta$. The determination

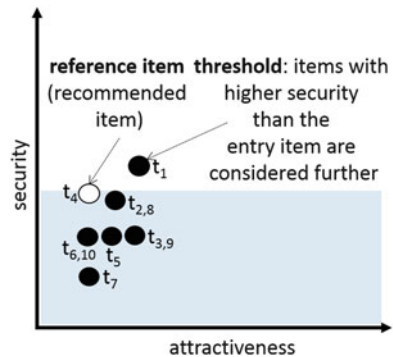
of the *complete set of minimal diagnoses* for a set of requirements inconsistent with the underlying item table (the corresponding conjunctive query results in \emptyset) is based on the construction of hitting set trees [72].

There are two possibilities of resolving the conflict set CS_1 . If we decide to delete the requirement r_3 , $\sigma_{(\{r_1, r_2\})}I \neq \emptyset$, i.e., a diagnosis has been identified ($\Delta_1 = \{r_3\}$) and—as a consequence—all CS_i have been resolved. Choosing the other alternative and resolving CS_1 by deleting r_1 does not result in a diagnosis since the conflict CS_2 is not resolved. Resolving CS_2 by deleting r_3 does not result in a minimal diagnosis, since r_3 already represents a diagnosis. The second (and last) minimal diagnosis that can be identified in our running example is $\Delta_2 = \{r_1, r_2\}$. For a detailed discussion of the underlying algorithm and analysis, we refer to [21, 72]. Note that a diagnosis provides a hint to which requirements have to be changed. For a discussion of how requirement repairs (changes) are calculated, we refer to Felfernig et al. [19].

Critiquing-Based Recommendation

Critiquing-based recommender systems support the navigation in the item space where in each critiquing cycle a reference item is presented to the user and the user either accepts the (recommended) item or searches for different solutions by specifying *critiques* (see Fig. 1.1). Critiquing-based recommender systems are useful in situations where users are not experts in the item domain and prefer to specify their requirements on the level of critiques [46]. Critiques are basic criteria that are used for determining new recommendations which take into account the (changed) preferences of the current user. Examples of such critiques in the context of our running example are *higher level of security* and *higher attractiveness*. Critiques are used as search criteria to identify corresponding candidate items, i.e., items that are shown to the user if he/she has specified critiques on the current reference item. Critiquing-based recommenders often search for items that are similar to the current reference item and additionally take into account the new

Fig. 1.1 Example of a critiquing scenario: an item (t_4) is shown as *reference item* to the user. The user specifies the critique “higher security.” The new reference item is t_1 since it is consistent with the critique and the item most similar to t_4 (here, it is also the only remaining alternative)



criteria defined as critiques. When searching for a new reference item, similarity and diversity metrics are applied to systematically guide the navigation in the item space [43].

Similarity metrics in critiquing-based approaches are used to determine the similarity between a reference item currently shown to the user and a set of candidate items to be shown to the user in the next critiquing cycle. Example similarity metrics often used in the context of critiquing-based recommendation scenarios are represented by Formulae 1.5–1.9. In this context, $sim(r, c)$ denotes the similarity between the reference item r and the candidate item c , the subroutine $s(r.i, c.i)$ is represented by different *attribute-level similarity metrics* (MIB, LIB, NIB, and EIB). If a higher attribute value is better than a lower one, *More-Is-Better* (MIB) (Formula 1.6) is used to evaluate the attribute of a candidate item ($c.i$). Vice versa, if low attributes values are considered better, the *Less-Is-Better* (LIB) similarity metric is used (Formula 1.7). Furthermore, *Nearer-Is-Better* (NIB) (Formula 1.8) is used if the attribute value of the candidate item should be as near as possible to the attribute $r.i$. Finally, *Equal-Is-Better* (EIB) is used in situations where attribute values should be equal (Formula 1.9). Besides taking into account the similarity between the reference item and candidate items, some critiquing-based systems also take into account the *compatibility* of a candidate item with regard to the complete critiquing history [61]. Thus, a trade-off between similarity to the reference item and critique compatibility can be achieved.

$$sim(r, c) = \sum_{i \in attributes} s(r.i, c.i) * w(i) \quad (\sum_{i \in attributes} w(i) = 1) \quad (1.5)$$

$$MIB : s(r.i, c.i) = \frac{val(c.i) - minval(r.i)}{maxval(r.i) - minval(r.i)} \quad (1.6)$$

$$LIB : s(r.i, c.i) = \frac{maxval(r.i) - val(c.i)}{maxval(r.i) - minval(r.i)} \quad (1.7)$$

$$NIB : s(r.i, c.i) = 1 - \frac{|val(r.i) - val(c.i)|}{maxval(r.i) - minval(r.i)} \quad (1.8)$$

$$EIB : s(r.i, c.i) = \begin{cases} 1 & \text{if } r.i = c.i \\ 0 & \text{otherwise} \end{cases} \quad (1.9)$$

If users are knowledgeable in the item domain, the application of search-based approaches such as constraint-based recommendation makes more sense. Different types of critiquing-based approaches primarily differ in terms of the way in which user preferences can be specified. *Unit critiquing* [10, 54] only supports the definition of critiques (change requests) that are related to a single item property (attribute). *Compound critiques* allow the specification of change requests over multiple item properties and thus allow to reduce the number of needed critiquing cycles [60]. Finally, *experience-based critiquing* takes into account critiquing

histories of previous users to better predict reference items and thus to reduce the number of needed interaction cycles [54, 62]. For an in-depth discussion of different (additional) variants of critiquing-based recommendation, we refer to [10, 11, 32, 54, 60, 73].

Hybrid Recommendation

After having discussed the basic recommendation approaches of *collaborative filtering*, *content-based filtering*, *constraint-based*, and *critiquing-based recommendation*, we will now present some possibilities to combine these.

A major motivation for the development of hybrid recommender systems is the opportunity to achieve a better accuracy [9]. There are different approaches to evaluate the accuracy of recommendation algorithms. These approaches can be categorized into *predictive accuracy metrics* such as the mean absolute error (MAE), *classification accuracy metrics* such as precision and recall, and *rank accuracy metrics* such as Kendall's Tau (see Chap. 3). For a discussion of accuracy metrics in recommendation scenarios for individual users, we refer to Gunawardana and Shani [34] and Jannach et al. [43].

We now take a look at *example design types* of hybrid recommenders [9, 43]. These are *weighted*, *mixed*, and *cascade* (see Table 1.11). The basic assumption in the following example is that individual recommendation approaches return a list of *five* recommended items where each item has an assigned (recommender-individual) prediction *out of* {1.0, 2.0, 3.0, 4.0, 5.0}. For a more detailed discussion of hybridization strategies, we refer the reader to Burke [9] and Jannach et al. [43].

Weighted Weighted hybrid recommendation is based on the idea of deriving recommendations by combining the results (predictions) computed by individual recommenders. An example thereof is depicted in Table 1.12 where the individual

Table 1.11 Examples of hybrid recommendation approaches (*RECS* = set of recommenders, s = recommender-individual prediction, *score* = item score)

Method	Description	Calculation
Weighted	Predictions (s) of individual recommenders are summed up	$\text{score}(\text{item}) = \sum_{\text{rec} \in \text{RECS}} s(\text{item}, \text{rec})$
Mixed	Recommender-individual predictions (s) are combined into one recommendation result	$\text{score}(\text{item}) = \text{zipper-function}(\text{item}, \text{RECS})$
Cascade	The prediction of one recommender ($n - 1$) is used as input for the next recommender (n)	$\text{score}(\text{item}) = s(\text{item}, \text{rec}_n) \leftarrow s(\text{item}, \text{rec}_{n-1}) \leftarrow \dots$

Table 1.12 Example of *weighted* hybrid recommendation: individual predictions are integrated into a score

Items	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}
$s(t_i, \text{collaborative filtering})$	1.0	3.0	—	5.0	—	2.0	—	4.0	—	—
$s(t_i, \text{content-based filtering})$	—	1.0	2.0	—	—	3.0	4.0	5.0	—	—
$\text{score}(t_i)$	1.0	4.0	2.0	5.0	0.0	5.0	4.0	9.0	0.0	0.0
$\text{ranking}(t_i)$	7	4	6	2	8	3	5	1	9	10

Item t_8 receives the best overall score (9.0)

Table 1.13 Example of *mixed* hybrid recommendation: individual predictions are integrated into one score conform the zipper principle (best collaborative filtering prediction receives score=10.0, best content-based filtering prediction receives score=9.0 and so forth)

Items	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}
$s(t_i, \text{collaborative filtering})$	1.0	3.0	—	5.0	—	2.0	—	4.0	—	—
$s(t_i, \text{content-based filtering})$	—	1.0	2.0	—	—	3.0	4.0	5.0	—	—
$\text{score}(t_i)$	4.0	8.0	5.0	10.0	0.0	6.0	7.0	9.0	0.0	0.0
$\text{ranking}(t_i)$	7	3	6	1	8	5	4	2	9	10

item scores of a collaborative and a content-based recommender are summed up. Item t_8 receives the highest overall score (9.0) and is ranked highest by the weighted hybrid recommender.⁶

Mixed Mixed hybrid recommendation is based on the idea that predictions of individual recommenders are shown in one integrated result. For example, the results of a collaborative filtering and a content-based recommender can be ranked as sketched in Table 1.13. Item scores can be determined, for example, on the basis of the zipper principle, i.e., the item with the highest collaborative filtering prediction value receives the highest overall score (10.0), the item with the best content-based filtering prediction value receives the second best overall score, and so forth.

Cascade The basic idea of cascade-based hybridization is that recommenders in a pipeline of recommenders exploit the recommendation of the upstream recommender as a basis for deriving their own recommendation. The constraint-based recommendation approach is an example of a cascade-based hybrid recommendation approach. First, items that are consistent with the given requirements are preselected by a conjunctive query Q . We can assume, for example, that $s(\text{item}, Q) = 1.0$ if the item has been selected and $s(\text{item}, Q) = 0.0$ if the item has not been selected. In our case, the set of requirements $\{r_1 : \text{topics} = \text{nature}\}$ in the running example leads to the selection of the items $\{t_2, t_4, t_7, t_8\}$. Thereafter, these items are ranked conform to their utility for the current user (utility-based

⁶If two or more items have the same overall score, a possibility is to force a decision by lot; where needed, this approach can also be applied by other hybrid recommendation approaches.

ranking U). Based on the entries of Table 1.10, a utility-based ranking (U) would determine the item order $\text{utility}(t_2) \geq \text{utility}(t_8) > \text{utility}(t_4) > \text{utility}(t_7)$, assuming that the current user assigns a weight of 0.6 to the interest dimension *security* ($\text{weight}(u_a, \text{security}) = 0.6$), a weight of 0.3 to the interest dimension *attractiveness* ($\text{weight}(u_a, \text{attractiveness}) = 0.3$), and a weight of 0.1 to the interest dimension *crowdedness* ($\text{weight}(u_a, \text{crowdedness}) = 0.1$). In this example, the recommender Q is the first one and the results of Q are forwarded to the utility-based recommender.

Further Approaches Further hybrid recommendation approaches are the following [9]. *Switching* denotes an approach where—depending on the current situation—a specific recommendation approach is chosen. For example, if a user has a low level of product knowledge, then a critiquing-based recommender will be chosen. Vice-versa, if the user is an expert, an interface will be provided where the user is enabled to explicitly state his/her preferences on a detailed level. *Feature combination* denotes an approach where different data sources are exploited by a single recommender. For example, a recommendation algorithm could exploit semantic item knowledge in combination with item ratings (see Table 1.8).

Outlook Due to the increasing popularity of social platforms and online communities, group recommender systems are becoming an increasingly important technology [36, 58]. Example application domains of group recommendation technologies include tourism [61] (e.g., *which hotels or travel destinations should be visited by a group?*) and interactive television [57] (*which sequence of television programs will be accepted by a group?*). For the most part, group recommendation algorithms are operating on *simple items* such as hotels, travel destinations, and television programs. Examples of *complex items* are cars, round trips, and software release plans. In the remainder of this book, we will discuss different approaches that determine item recommendations for groups.

1.4 Relationship Between Algorithms and Choice Patterns

As already mentioned, recommender systems can be regarded as important support for human decision making [42]. In the following, we will briefly discuss different patterns of choice, i.e., approaches people use to solve a decision task/problem. Following the concepts introduced in [41, 42], we will explain these patterns and point out related recommendation approaches. Our major motivation for this discussion is the increasingly perceived relevance of human decision making aspects in recommendation contexts [42, 55].

Socially-Influenced Choice If someone is interested in purchasing a digital camera, has no experiences in the item domain, but knows friends who are photography enthusiasts, there is a high probability that the opinion of these friends can have an impact on the camera purchase decision. People are influenced by the opinions and advice of friends beyond item recommendation, for example, in terms of social

expectations of what is “cool” or what is “politically correct” [42]. From the viewpoint of recommender systems, *collaborative filtering* [48] simulates socially-influenced decisions where the preferences of nearest neighbors (social environment of the user) are taken as input for recommendations proposed to the current user. In such recommendation scenarios, *trust* plays an important role since only the preferences of trusted users should be taken into account by the recommendation algorithm [30, 65]. If available, an important aspect to be taken into account are *social networks* which can serve as a basis for representing a user’s trust networks and also to determine the nearest neighbors relevant in the recommendation context [35]. In scenarios where groups of users have to make a decision, the personal opinion of a group member is in many cases influenced by his/her surroundings, i.e., other group members or even other groups [71]. Consequently, a key feature of group decisions is that social influence occurs not only in the environment but also within the group of decision makers themselves (see Chap. 9)—a fact that makes the socially-based pattern even more important than it is for individual choices.

Attribute-Based Choice This pattern is based on the idea that individual alternatives (items) are described by attributes which can be associated with importance values that reflect individual relevance. People can in principle evaluate an item by considering the values and importances of its attributes and computing something like a weighted sum for each item. More typically, they apply less effortful strategies. For example, *attributes* can serve as criteria to figure out whether an alternative should be taken into account (also denoted as *elimination by aspects* where users are kept in the loop), serve as a basis for the ranking of the available alternatives (utility-based ranking), or can be used for both purposes [5, 42, 67]. Due to the fact that importance of attributes and preferences for particular attribute levels are not stable and are not known in detail from the beginning [4], decision processes are iterative where users change or slightly adapt their preferences. Recommenders often integrated in such scenarios are *utility-based* and *constraint-based recommenders* [8, 14]. Constraint-based recommendation [14] is regarded as a special type of knowledge-based recommendation approach where the recommendation knowledge is represented in terms of explicitly defined attributes and constraints. The second type of knowledge-based recommendation (if we interpret item attributes and corresponding critiques as semantic knowledge) is represented by critiquing-based systems [11] where users specify change requests with regard to a reference alternative and the recommender system proposes a new candidate alternative that takes into account previous critiques. Critiques are typically defined on item attributes, therefore critiquing-based approaches can also be helpful to support attribute-based decision processes. There are situations where attribute-based decisions have to be supported for groups. Group members have to develop consensus with regard to a set of attributes, for example, arriving at a common set of importance weights or agreeing on the selection of an attribute value. If contradicting preferences occur, visualization and analysis methods can be applied to resolve inconsistencies [26].

Trial & Error Based Choice This choice pattern is often used in situations where the item domain is unknown and users want to better understand the domain and

analyze the pros and cons of the different existing options in more detail [41]. For example, customers purchase the trial version of a software, test different cars, or evaluate some potential holidays within the scope of a short-term visit. In the context of recommender systems, trial & error based decisions are supported, for example, by critiquing-based recommendation approaches [11] where a user analyzes individual recommendations and then defines criteria that should additionally be taken into account when presenting the next item. Critiquing-based approaches support explorative navigation in large search spaces which better helps to develop an understanding of the item domain [43]. Critiquing-based recommendation also plays a role when it comes to determining recommendations for groups. For example, a group makes a decision with regard to a skiing resort to visit in the next winter season [61].

Experience-Based Choice This type of choosing can be applied to situations where experiences of users from the past directly influence current decisions, often making it unnecessary to apply any of the other choice patterns (see, e.g., [3]). For example, if someone purchased a book from a specific author or a specific publisher, this experience can be exploited in future purchasing scenarios and books from one consideration set could be preferred based on these experience-based criteria. Another example is car purchasing: if a customer is completely satisfied with his/her car (specific brand), there is a good chance that he/she will stick with the brand when purchasing a new car. In this context, positive feelings from the past trigger positive feelings about specific alternatives in ongoing decision processes [42]. From the point of view of recommendation approaches, content-based recommendation implements experience-based decisions in the sense that positively-rated purchases in the past represent major criteria to recommend similar items in the future. For example, a user liked a book of a specific author, therefore a new book from the same author is recommended to the user. We want to point out that in the recommender systems community, critiquing-based recommendation [11] is *also* considered as specific type of case-based recommendation [53, 77] where items are regarded as cases and items similar to a given set of user requirements (the case description) are recommended. Since cases represent experiences from the past, critiquing-based recommendation can also be considered as representative of experience-based decisions. This becomes even clearer if we take a look at recent experience-based critiquing approaches where experiences from previous recommendation sessions are taken into account and users with similar critiquing histories receive similar recommendations [54, 62]. Here, the experience-based pattern is supplemented with the socially-based, since the experiences of other persons are being exploited as well.

Consequence-Based Choice An alternative decision strategy is to think about the potential consequences of choosing a specific alternative. In contrast to decision patterns such as attribute-based decisions and socially-influenced decisions, thinking about the consequences of making a specific decision (choosing a specific alternative) constitutes an additional dimension of a decision process. One challenge in this context is the uncertainty about the consequences of taking a decision

and how to best deal with this. Since anticipating and evaluating consequences is typically an effortful process, people are more likely to do so in some domains than in others. Low involvement items, i.e., items with lower risk impacts related to suboptimal decisions (e.g., movies, restaurants, and songs) will result in less investment in the analysis of the impact of choosing such items. In contrast, high-involvement items, i.e., items with higher associated risks, will result in higher investment in the analysis of the consequences of choosing a specific alternative. We want to point out that consequence-based decision patterns are often orthogonal, i.e., they play a role in combination with each of the mentioned recommendation approaches. Like all of the choice patterns, the consequence-based pattern can be combined with other patterns. For example, a chooser may apply the attribute-based pattern to quickly form a consideration set and then apply the consequence-based pattern to the consideration set. An important aspect of decision making is the relatively strong emphasis on the potential negative consequences of a decision—this aspect is taken into account in *prospect theory* [45], an asymmetric utility function which postulates that losses have a higher negative evaluation compared to equal gains.

Policy-Based Choice In the initial phase of a decision process, it can be the case that decision makers also define their decision policy. For example, during the weekend, if a movie recommender proposes a thriller and a cartoon, a family may apply a general policy of first consuming the movie that is suitable for children. They don't need to think in each case about the justification for this policy (i.e., that children need to go to bed earlier). When deciding about the next restaurant for a Christmas party, a company could have predefined the (meta-level) decision policy of *majority voting*, i.e., the majority defines which restaurant will be chosen for the next Christmas party. If two books with equal average ratings and high-quality reviews are in a consideration set, the customer could purchase the book from a publisher he/she previously found satisfying. A book customer may (for any of various reasons) have acquired the policy of buying a particular type of book from a particular publisher. In that case, the customer does not even need to consider books from other publishers. As these examples suggest, the policy-based decision pattern can play a role in the context of each basic recommendation scenario (*collaborative*, *content-based*, *constraint-based*, and *critiquing-based*).

1.5 Book Overview

In this chapter, we introduced recommender systems as a basic technology to support different decision scenarios. We gave an overview of recommendation algorithms and showed their application in the context of a scenario from the travel domain. Finally, we discussed the relationship between basic patterns of human choice and related supportive recommendation approaches. *The remainder of this book is organized as follows.* In Chap. 2, we provide an overview of group

recommendation algorithms with the goal of showing how basic recommendation algorithms (collaborative filtering, content-based filtering, constraint-based, and critiquing-based recommendation) can be tailored to group settings. In this context, we show the relationship of group recommenders to the algorithms discussed in Chap. 1. In Chap. 3, we sketch approaches to evaluate group recommender systems. An overview of existing applications of group recommendation technologies is provided in Chap. 4. In Chap. 5, we focus on a discussion of ways to elicit and manage user preferences. Chapter 6 deals with explanation approaches in the context of group recommendation. In Chap. 7, we introduce additional decision scenarios, for example, group-based configuration, group-based resource balancing, and group-based release planning. Chapter 8 analyzes the existence and ways to counteract decision biases that can occur in the context of group decision making. Chapter 9 focuses on the role of personality and emotion in group recommendation. Finally, the book is concluded with a summary and an outlook (see Chap. 10).

References

1. C. Aggarwal, *Recommender Systems: The Textbook* (Springer, New York, 2016)
2. L. Ardissono, A. Felfernig, G. Friedrich, A. Goy, D. Jannach, G. Petrone, R. Schäfer, M. Zanker, A framework for the development of personalized, distributed web-based configuration systems. *AI Mag.* **24**(3), 93–108 (2003)
3. T. Betsch, S. Haberstroh, *The Routines of Decision Making* (Lawrence Erlbaum Associates, Mahwah, NJ, 2005)
4. J. Bettman, M. Luce, J. Payne, Constructive consumer choice processes. *J. Consum. Res.* **25**(3), 187–217 (1998)
5. S. Bhatia, Associations and the accumulation of preference. *Psychol. Rev.* **120**(3), 522–543 (2013)
6. D. Billsus, M. Pazzani, Learning collaborative information filters, in *15th International Conference on Machine Learning (ICML'98)*, pp. 46–54 (1998)
7. L. Boratto, S. Carta, State-of-the-art in group recommendation and new approaches for automatic identification of groups, in *Information Retrieval and Mining in Distributed Environments*. Studies in Computational Intelligence, vol. 324 (Springer, Heidelberg, 2011), pp. 1–20
8. R. Burke, Knowledge-based recommender systems. *Encycl. Library Inform. Syst.* **69**(32), 180–200 (2000)
9. R. Burke, Hybrid recommender systems: survey and experiments. *User Model. User-Adap. Inter. (UMUAI)* **12**(4), 331–370 (2002)
10. R. Burke, K. Hammond, B. Young, The FindMe approach to assisted browsing. *IEEE Expert: Intell. Syst. Appl.* **12**(4), 32–40 (1997)
11. L. Chen, P. Pu, Critiquing-based recommenders: survey and emerging trends. *User Model. User-Adap. Inter. (UMUAI)* **22**(1–2), 125–150 (2012)
12. M. Ekstrand, J. Riedl, J. Konstan, Collaborative filtering recommender systems. *Found. Trends Human-Comput. Interact.* **4**(2), 81–173 (2011)
13. A. Falkner, A. Felfernig, A. Haag, Recommendation technologies for configurable products. *AI Mag.* **32**(3), 99–108 (2011)
14. A. Felfernig, R. Burke, Constraint-based recommender systems: technologies and research issues, in *ACM International Conference on Electronic Commerce (ICEC08)*, Innsbruck, Austria, pp. 17–26 (2008)

15. A. Felfernig, G. Friedrich, D. Jannach, M. Stumptner, Consistency-based diagnosis of configuration knowledge bases. *Artif. Intell.* **152**(2), 213–234 (2004)
16. A. Felfernig, G. Friedrich, D. Jannach, M. Zanker, An integrated environment for the development of knowledge-based recommender applications. *Int. J. Electron. Commer. (IJEC)* **11**(2), 11–34 (2006)
17. A. Felfernig, G. Friedrich, L. Schmidt-Thieme, Guest editors' introduction: recommender systems. *IEEE Intell. Syst.* **22**(3), 18–21 (2007)
18. A. Felfernig, B. Gula, G. Leitner, M. Maier, R. Melcher, E. Teppan, Persuasion in knowledge-based recommendation, in *3rd International Conference on Persuasive Technology*. Lecture Notes in Computer Science (Springer, Berlin, 2008), pp. 71–82
19. A. Felfernig, G. Friedrich, M. Schubert, M. Mandl, M. Mairitsch, E. Teppan, Plausible repairs for inconsistent requirements, in *Proceedings of the IJCAI'09*, Pasadena, CA, pp. 791–796 (2009)
20. A. Felfernig, W. Maalej, M. Mandl, M. Schubert, F. Ricci, Recommendation and decision technologies for requirements engineering, in *ICSE 2010 Workshop on Recommender Systems in Software Engineering (RSSE 2010)*, Cape Town, pp. 11–15 (2010)
21. A. Felfernig, M. Schubert, C. Zehentner, An efficient diagnosis algorithm for inconsistent constraint sets. *Artif. Intell. Eng. Des. Anal. Manuf. (AIEDAM)* **26**(1), 53–62 (2012)
22. A. Felfernig, M. Schubert, S. Reiterer, Personalized diagnosis for over-constrained problems, in *23rd International Conference on Artificial Intelligence (IJCAI 2013)*, Peking, China, pp. 1990–1996 (2013)
23. A. Felfernig, M. Jeran, G. Ninaus, F. Reinfrank, S. Reiterer, Toward the next generation of recommender systems, in *Multimedia Services in Intelligent Environments: Recommendation Services* (Springer, Heidelberg, 2013), pp. 81–98
24. A. Felfernig, M. Jeran, G. Ninaus, F. Reinfrank, S. Reiterer, M. Stettinger, Basic approaches in recommendation systems, in *Recommendation Systems in Software Engineering* (Springer, Berlin, 2013), pp. 15–37
25. A. Felfernig, L. Hotz, C. Bagley, J. Tiihonen, *Knowledge-Based Configuration: From Research to Business Cases*, 1st edn. (Elsevier/Morgan Kaufmann, Waltham, MA, 2014)
26. A. Felfernig, M. Atas, T.N. Trang Tran, M. Stettinger, Towards group-based configuration, in *International Workshop on Configuration 2016 (ConfWS'16)*, pp. 69–72 (2016)
27. 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, pp. 335–344 (2017)
28. H. Garcia-Molina, G. Koutrika, A. Parameswaran, Information seeking: convergence of search, recommendations, and advertising. *Commun. ACM* **54**(11), 121–130 (2011)
29. M. Gasparic, A. Janes, What recommendation systems for software engineering recommend: a systematic literature review. *J. Syst. Softw.* **113**, 101–113 (2016)
30. J. Golbeck, *Computing with Social Trust* (Springer, London, 2009)
31. D. Goldberg, D. Nichols, B. Oki, D. Terry, Using collaborative filtering to weave an information tapestry. *Commun. ACM* **35**(12), 61–70 (1992)
32. P. Gräsch, 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
33. R. Grünig, R. Kühn, *Successful Decision-Making* (Springer, Heidelberg, 2013)
34. A. Gunawardana, G. Shani, A survey of accuracy evaluation metrics of recommendation tasks. *J. Mach. Learn. Res.* **10**, 2935–2962 (2009)
35. J. He, W. Chu, A social network-based recommender system (SNRS), in *Data Mining for Social Network Data*. Annals of Information Systems, vol. 12 (Springer, New York, 2010), pp. 47–74
36. T. Hennig-Thurau, A. Marchand, P. Marx, Can automated group recommender systems help consumers make better choices? *J. Market.* **76**(5), 89–109 (2012)

37. J. Herlocker, J. Konstan, A. Borchers, J. Riedl, An algorithmic framework for performing collaborative filtering, in *Conference on Research and Development in Information Retrieval*, Berkeley, CA, pp. 230–237 (1999)
38. C. Huffman, B. Kahn, Variety for sale: mass customization or mass confusion? *J. Retail.* **74**(4), 491–513 (1998)
39. A. Jameson, More than the sum of its members: challenges for group recommender systems, in *International Working Conference on Advanced Visual Interfaces*, pp. 48–54 (2004)
40. 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, Heidelberg, 2007), pp. 596–627
41. A. Jameson, B. Berendt, S. Gabrielli, F. Cena, C. Gena, F. Vernero, K. Reinecke, Choice architecture for human-computer interaction. *Foundations and Trends in Human-Computer Interaction*, vol. 7 (Now Publishers Inc., Hanover, MA, 2014)
42. 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, New York, 2015), pp. 611–648
43. D. Jannach, M. Zanker, A. Felfernig, G. Friedrich, *Recommender Systems – An Introduction* (Cambridge University Press, New York, 2010)
44. U. Junker, QUICKXPLAIN: preferred explanations and relaxations for over-constrained problems, in *19th International Conference on Artificial Intelligence, AAAI'04* (AAAI Press, San Jose, 2004), pp. 167–172
45. D. Kahneman, A. Tversky, Prospect theory: an analysis of decision under risk. *Econometrica* **47**(2), 263–291 (1979)
46. B. Knijnenburg, N. Reijmer, M. Willemsen, Each to his own: how different users call for different interaction methods in recommender systems, in *RecSys 2011*, Chicago, IL, pp. 141–148 (2011)
47. M. Kompan, M. Bielikova, Group recommendations: survey and perspectives. *Comput. Inform.* **33**(2), 446–476 (2014)
48. J. Konstan, B. Miller, D. Maltz, J. Herlocker, L. Gordon, J. Riedl, GroupLens: applying collaborative filtering to usenet news. *Commun. ACM* **40**(3), 77–87 (1997)
49. Y. Koren, R. Bell, C. Volinsky, Matrix factorization techniques for recommender systems. *IEEE Comput.* **42**(8), 30–37 (2009)
50. G. Leitner, A. Fercher, A. Felfernig, K. Isak, S. Polat Erdeniz, A. Akcay, M. Jeran, Recommending and configuring smart home installations, in *International Workshop on Configuration 2016 (ConfWS'16)*, pp. 17–22 (2016)
51. J. Levin, B. Nalebuff, An introduction to vote-counting schemes. *J. Econ. Perspect.* **9**(1), 3–26 (1995)
52. G. Linden, B. Smith, J. York, Amazon.com recommendations – item-to-item collaborative filtering. *IEEE Internet Comput.* **7**(1), 76–80 (2003)
53. F. Lorenzi, F. Ricci, Case-based recommender systems: a unifying view, in *International Conference on Intelligent Techniques for Web Personalization*, pp. 89–113 (2003)
54. M. Mandl, A. Felfernig, Improving the performance of unit critiquing, in *20th International Conference on User Modeling, Adaptation, and Personalization (UMAP 2012)*, Montreal, pp. 176–187 (2012)
55. M. Mandl, A. Felfernig, E. Teppan, M. Schubert, Consumer decision making in knowledge-based recommendation. *J. Intell. Inform. Syst.* **37**(1), 1–22 (2010)
56. A. Marchand, *Empfehlungssysteme für Gruppen* (EUL Verlag, 2011)
57. J. Masthoff, Group modeling: selecting a sequence of television items to suit a group of viewers. *User Model. User-Adap. Inter. (UMUAI)* **14**(1), 37–85 (2004)
58. J. Masthoff, Group recommender systems: combining individual models, in *Recommender Systems Handbook* (Springer, New York, 2011), pp. 677–702
59. J. McCarthy, T. Anagnost, MusicFX: an arbiter of group preferences for computer supported collaborative workouts, in *Conference on Computer Support Cooperative Work*, Seattle, WA, pp. 363–372 (1998)

60. 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
61. 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)
62. K. McCarthy, Y. Salem, B. Smyth, Experience-based critiquing: reusing critiquing experiences to improve conversational recommendation, in *International Conference on Case-Based Reasoning (ICCBR 2010)*, Alessandria, pp. 480–494 (2010)
63. H. Mengash, A. Brodsky, A group recommender for investment in microgrid renewable energy sources, in *50th Hawaii International Conference on System Sciences*, Hawaii, pp. 1485–1494 (2017)
64. G. Ninaus, A. Felfernig, M. Stettinger, S. Reiterer, G. Leitner, L. Weninger, W. Schanil, INTELLIREQ: Intelligent techniques for software requirements engineering, in *Prestigious Applications of Intelligent Systems Conference (PAIS)*, pp. 1161–1166 (2014)
65. J. O'Donovan, B. Smyth, Trust in recommender systems, in *ACM UII 2005*, San Diego, CA, pp. 167–174 (2005)
66. D. Paraschakis, Recommender systems from an industrial and ethical perspective, in *10th ACM Conference on Recommender Systems*, Boston, MA, pp. 463–466 (2016)
67. J. Payne, J. Bettman, E. Johnson, *The Adaptive Decision Maker* (Cambridge University Press, New York, 1993)
68. M. Pazzani, D. Billsus, Learning and revising user profiles: the identification of interesting web sites. *Mach. Learn.* **27**(3), 313–331 (1997)
69. K. Peniwati, Criteria for evaluating group decision-making methods. *Math. Comput. Model.* **46**(7–8), 935–947 (2007)
70. R. Petty, J. Cacioppo, D. Schumann, Central and peripheral routes to advertising effectiveness: the moderating role of involvement. *J. Consum. Res.* **10**(2), 135–146 (1983)
71. L. Recalde, A social framework for set recommendation in group recommender systems, in *European Conference on Information Retrieval* (Springer, New York, 2017), pp. 735–743
72. R. Reiter, A theory of diagnosis from first principles. *Artif. Intell. J.* **32**(1), 57–95 (1987)
73. F. Ricci, Q. Nguyen, Acquiring and revising preferences in a critique-based mobile recommender systems. *IEEE Intell. Syst.* **22**(3), 22–29 (2007)
74. B. Sarwar, G. Karypis, J. Konstan, J. Riedl, Item-based collaborative filtering recommendation algorithms, in *10th WWW Conference*, pp. 285–295 (2001)
75. G. Sielis, C. Mettouris, G. Papadopoulos, A. Tzanavari, R. Dols, Q. Siebers, A context aware recommender system for creativity support tools. *J. Universal Comput. Sci.* **17**(12), 1743–1763 (2011)
76. B. Smith, G. Linden, Two decades of recommender systems at Amazon.com. *IEEE Internet Comput.* **21**(3), 12–18 (2017)
77. B. Smyth, Case-based recommendation, in *The Adaptive Web*, ed. by P. Brusilovsky, A. Kobsa, W. Nejdl. *Lecture Notes in Computer Science*, vol. 4321 (Springer, Berlin/Heidelberg, 2007), pp. 342–376
78. M. Stettinger, A. Felfernig, CHOICLA: Intelligent decision support for groups of users in context of personnel decisions, in *ACM RecSys'2014 IntRS Workshop*, Foster City, CA, pp. 28–32 (2014)
79. 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, GA, pp. 148–157 (2015)
80. N. Tintarev, J. Masthoff, Designing and evaluating explanations for recommender systems, in *Recommender Systems Handbook* (Springer, New York, 2011), pp. 479–510
81. N. Tintarev, J. O'Donovan, A. Felfernig, Human interaction with artificial advice givers. *ACM Trans. Interact. Intell. Syst.* **6**(4), 1–10 (2016)
82. A. Valdez, M. Zieffle, K. Verbert, A. Felfernig, A. Holzinger, Recommender systems for health informatics: state-of-the-art and future perspectives, in *Machine Learning for Health Informatics* (Springer, Cham, 2016), pp. 391–414