

Chapter 1

Recommender Systems: Introduction and Challenges

Francesco Ricci, Lior Rokach, and Bracha Shapira

1.1 Introduction

Recommender Systems (RSs) are software tools and techniques that provide suggestions for items that are most likely of interest to a particular user [17, 41, 42]. The suggestions relate to various decision-making processes, such as what items to buy, what music to listen to, or what online news to read.

“Item” is the general term used to denote what the system recommends to users. An RS normally focuses on a specific type of item (e.g., CDs or news) and, accordingly its design, its graphical user interface, and the core recommendation technique used to generate the recommendations are all customized to provide useful and effective suggestions for that specific type of item.

RSs are primarily directed toward individuals who lack the sufficient personal experience or competence in order to evaluate the potentially overwhelming number of alternative items that a website, for example, may offer [42]. A prime example is a book recommender system that assists users in selecting a book to read. On the popular website, Amazon.com, the site employs an RS to personalize the online store for each customer [32]. Since recommendations are usually personalized, different users or user groups benefit from diverse, tailored suggestions. In addition, there are also non-personalized recommendations. These are much simpler to generate and are normally featured in magazines or newspapers. Typical examples

F. Ricci (✉)

Faculty of Computer Science, Free University of Bozen-Bolzano, Bolzano, Italy
e-mail: fricci@unibz.it

L. Rokach • B. Shapira

Department of Information Systems Engineering,
Ben-Gurion University of the Negev, Beer-Sheva, Israel
e-mail: liorrk@bgu.ac.il; bshapira@bgu.ac.il

include the top ten selections of books, CDs etc. While they may be useful and effective in certain situations, these types of non-personalized recommendations are not typically addressed by RS research.

In their simplest form, personalized recommendations are offered as ranked lists of items. In performing this ranking, RSs try to predict what the most suitable products or services are, based on the user's preferences and constraints. In order to complete such a computational task, RSs collect information from users regarding their preferences which are either explicitly expressed, e.g., as ratings for products or are inferred by interpreting the actions of the user. For instance, an RS may consider the navigation to a particular product page as an implicit sign of preference for the items shown on that page.

The development of RSs initiated from a rather simple observation: individuals often rely on recommendations provided by others in making routine, daily decisions [41, 51]. For example, it is common to rely on what one's peers recommend when selecting a book to read; employers count on recommendation letters in their recruiting decisions; and when selecting a movie to watch, individuals tend to read and rely on the movie reviews that a film critic has written, which appear in the newspaper they read.

In seeking to mimic this behavior, the first RSs applied algorithms in order to leverage recommendations produced by a community of users and deliver these recommendations to an "active" user, or a user looking for suggestions. The recommendations were for items that similar users, or those with similar tastes, had liked. This approach is termed collaborative-filtering and its rationale follows that if the active user agreed in the past with certain users, then the other recommendations coming from these similar users should be relevant as well as of interest to the active user.

As e-commerce websites began to develop, a pressing need emerged for providing recommendations derived from filtering the whole range of available alternatives. Users found it difficult to arrive at the most appropriate choices from the immense variety of items (products and services) that these websites offered.

The explosive growth and variety of information available on the Web and the rapid introduction of new e-business services (selling products, product comparison, auctions, etc.) frequently overwhelmed users, leading them to make poor decisions. The availability of choices, instead of producing a benefit, started to decrease users' well-being. It was understood that while choice is good, more choice is not always better. Indeed, choice, with its implications of freedom, autonomy, and self-determination can become excessive, and ultimately create a sense that freedom may come to be regarded as a kind of misery-inducing tyranny [49].

In recent years, RSs have proven to be a valuable means of coping with the information overload problem. Ultimately an RS addresses this phenomenon by pointing a user toward new, not-yet-experienced items that may be relevant to the user's current task. Upon a user's request, which can be articulated depending on the recommendation approach by the user's context and need, RSs generate recommendations using various types of knowledge and data about users, the available items, and previous transactions stored in customized databases. The user

can then browse the recommendations. One may accept them or not and may provide, immediately or at a later stage, implicit or explicit feedback. This user action and feedback can be stored in the recommender database and may be used for generating new recommendations in the coming user-system interactions.

As previously noted, the study of recommender systems is relatively new compared to research in other classical information system tools and techniques (e.g., databases or search engines). Recommender systems emerged as an independent research area in the mid-1990s [7, 24, 41, 51]. In recent years, the interest in recommender systems has dramatically increased, as the following facts indicate:

1. Recommender systems play an important role in highly-rated Internet sites such as Amazon.com, YouTube, Netflix, Spotify, LinkedIn, Facebook, Tripadvisor, Last.fm, and IMDb. Moreover many media companies are now developing and deploying RSs as part of the services they provide to their subscribers. For example, Netflix, the online provider of on-demand streaming media, awarded a million dollar prize to the team that first succeeded in substantially improving the performance of its recommender system [31].
2. There are conferences and workshops dedicated specifically to the field, namely the Association of Computing Machinery's (ACM) Conference Series on Recommender Systems (RecSys), established in 2007. This conference stands as the premier annual event in recommender technology research and applications. In addition, sessions dedicated to RSs are frequently included in more traditional conferences in the area of databases, information systems and adaptive systems. Additional noteworthy conferences within this scope include: ACM's Special Interest Group on Information Retrieval (SIGIR); User Modeling, Adaptation and Personalization (UMAP); Intelligent User Interfaces (IUI); World Wide Web (WWW); and ACM's Special Interest Group on Management Of Data (SIGMOD).
3. At institutions of higher education around the world, undergraduate and graduate courses are now dedicated entirely to RSs, tutorials on RSs are very popular at computer science conferences, and a book introducing RSs techniques has been published as well [27]. Springer is publishing several books on specific topics in recommender systems in its series: Springer Briefs in Electrical and Computer Engineering. A large, new collection of articles dedicated to recommender systems applications to software engineering has also recently been published [46].
4. There have been several special issues in academic journals which cover research and developments in the RSs field. Among the journals that have dedicated issues to RSs are: AI Communications (2008); IEEE Intelligent Systems (2007); International Journal of Electronic Commerce (2006); International Journal of Computer Science and Applications (2006); ACM Transactions on Computer Human Interaction (2005); ACM Transactions on Information Systems (2004); User Modeling and User-Adapted Interaction (2014, 2012); ACM Transactions on Interactive Intelligent Systems (2013); and ACM Transactions on Intelligent Systems and Technology (2015).

In this introductory chapter, we briefly discuss basic RS ideas and concepts. Our main goal is not to present a self-contained comprehensive survey on RSs but rather to delineate, in a coherent and structured way, the chapters included in this handbook and to help the reader navigate the rich and detailed content that the handbook offers. The reader can also consult these recent introductions or surveys on recommender systems [13, 30, 34, 40, 44]. At the end of this chapter, we have identified some research challenges that we believe are particularly important for the future of the area.

The handbook is divided into five sections: recommendation techniques; recommender systems evaluation; recommender systems applications; recommender systems and human computer interaction; and advanced algorithms.

The first section presents the techniques most popularly used today for building RSs, such as collaborative filtering; content-based, data mining methods; and context-aware methods.

The second section surveys techniques and approaches that have been utilized to evaluate the quality of the recommendations. The section also considers aspects that may affect RS design (domain, device, interfaces, users, etc.). Finally, it discusses methods, challenges and measures to be applied in evaluating the developed systems with user experiments.

The third section includes papers dealing with a number of issues related to how recommendations are presented, browsed, explained and visualized. Among them, this section focuses on user's privacy and the decision making process supported by a recommender system.

The fourth section is fully dedicated to applications of recommender systems. We offer here a broad spectrum of the usage of these techniques in music, mobile computing, dating, social networks, education, and movies.

The last section presents papers on various advanced topics, such as: the exploitation of active learning principles to guide the acquisition of new knowledge; novelty and diversity in the recommendations; suitable techniques for protecting a recommender system against attacks of malicious users; and RSs that aggregate multiple types of user feedback and preferences to build more reliable recommendations.

1.2 Recommender Systems' Function

In the previous section, we defined RSs as software tools and techniques that provide users with suggestions for items that a user may wish to utilize. Now we wish to refine this definition to illustrate a range of possible roles that an RS can play. Firstly, we must distinguish between the role played by the RS on behalf of the service provider, from that of the user of the RS. For instance, a travel recommender system is typically introduced by a travel intermediary such as Expedia.com, or a destination management organization, such as Visitfinland.com, in order to increase its turnover or sell more hotel rooms in the case of Expedia, and increase the number

of tourists to the destination in the case of the destination management organization [14, 43]. The user's primary motivations for accessing the two systems would be to find a suitable hotel and interesting events or attractions when visiting a destination.

In fact, there are various reasons as to why service providers may want to exploit this technology:

- *Increase the number of items sold.* This is probably the most important function for a commercial RS, i.e., to be able to sell an additional set of items compared to those usually sold without any kind of recommendation. This goal is achieved because the recommended items are likely to suit the user's needs and wants. Presumably the user will recognize this after having tried several recommendations.¹ Non-commercial applications have similar goals, even if there is no cost for the user that is associated with selecting an item. For instance, a content network aims at increasing the number of news items read on its site. In general, we can say that from the service provider's point of view, the primary goal for introducing an RS is to increase the conversion rate, i.e., the number of users that accept the recommendation and consume an item, compared to the number of simple visitors that just browse through the information.
- *Sell more diverse items.* Another major function of an RS is to enable the user to select items that might be hard to find without a precise recommendation. For instance, in a tourist RS the service provider is interested in promoting all the places of interest in a tourist area, not just the most popular ones. This could be difficult without an RS since the service provider cannot afford the risk of advertising places that are not likely to suit a particular user's taste. Therefore, an RS suggests or advertises unpopular places to the right users.
- *Increase the user satisfaction.* A well designed RS can also improve the experience of the user with the site or the application. The user will find the recommendations interesting, relevant and, with a properly designed human-computer interaction, he or she will also enjoy using the system. The combination of effective, accurate recommendations and a usable interface will increase the user's subjective evaluation of the system. This, in turn, will increase system usage and the likelihood that the recommendations will be accepted.
- *Increase user fidelity.* A user should be loyal to a website which, when visited, recognizes the old customer and treats him as a valued visitor. This is a standard feature of an RS since many RSs compute recommendations, thus leveraging the information acquired from the user during previous interactions such as the user's ratings of items. Consequently, the longer the user interacts with the site, the more refined the user's model becomes: the system's representation of the user's preferences develops and the effectiveness of the recommender output to customize and match to the user's preferences is increased.

¹This issue, convincing the user to accept a recommendation, is discussed again when we explain the difference between predicting the user interest in an item and the likelihood that the user will select the recommended item.

- *Better understanding of what the user wants.* Another important function of an RS which can be leveraged to many other applications is the description of the user's preferences, which are collected either explicitly or predicted by the system. The service provider may then decide to reuse this knowledge for a number of other goals, such as improving the management of the item's stock or production. For instance, in the travel domain, destination management organizations can decide to advertise a specific region to new customer sectors or advertise a particular type of promotional message derived by analyzing the data collected by the RS (transactions of the users).

We mentioned above some important motivations as to why e-service providers introduce RSs. But users also may want an RS if it will effectively support their tasks or goals. Consequently an RS must balance the needs of these two players and offer a service that is valuable to both.

Herlocker et al. [26], in a paper that has become a classical reference in this field, define eleven popular tasks that an RS can assist in implementing. Some may be considered as the main or core tasks that are normally associated with an RS, such as offering suggestions for items that may be useful to a user. Others might be considered as more "opportunistic" ways to exploit an RS. As a matter of fact, this task differentiation is very similar to what happens with a search engine. Its primary function is to locate documents that are relevant to the user's information need, but it can also be used to check the importance of a webpage (looking at the position of the page in the result list of a query) or to discover the various usages of a word in a collection of documents.

- *Find Some Good Items:* Recommend to a user some items as a ranked list along with predictions of how much the user would like them (e.g., on a scale of one-to-five stars). This is the main recommendation task that many commercial systems address (see, for instance, Chap. 11). Some systems do not show the predicted rating.
- *Find all good items:* Recommend all the items that can satisfy some user needs. In such cases it is insufficient to just find some good items. This is especially true when the number of items is relatively small or when the RS is mission-critical, such as in medical or financial applications. In these situations, in addition to the benefit derived from carefully examining all the possibilities, the user may also benefit from the RS ranking of these items or from additional explanations that the RS generates.
- *Annotation in context:* Given an existing context, e.g., a list of items, emphasize some of them depending on the user's long-term preferences. For example, a TV recommender system might annotate which TV shows displayed in the electronic program guide (EPG) are worth watching (Chap. 15 provides interesting examples of this task).
- *Recommend a sequence:* Instead of focusing on the generation of a single recommendation, the idea is to recommend a sequence of items that is pleasing as a whole. Typical examples include recommending a TV series, a book on RSs after having recommended a book on data mining, or a compilation of musical tracks [28].

- *Recommend a bundle*: Suggest a group of items that fits well together. For instance, a travel plan may be composed of various attractions, destinations, and accommodation services that are located in a delimited area. From the point of view of the user, these various alternatives can be considered and selected as a single travel destination [45].
- *Just browsing*: In this task, the user browses the catalog without any imminent intention of purchasing an item. The task of the recommender is to help the user to browse the items that are more likely to fall within the scope of the user's interests for that specific browsing session. This is a task that has also been supported by adaptive hypermedia techniques [16].
- *Find credible recommender*: Some users do not trust recommender systems, thus they play with them to see how good they are at making recommendations. Hence, a certain system may also offer specific functions to let the users test its behavior in addition to those just required for obtaining recommendations.
- *Improve the profile*: This relates to the capability of the user to provide (input) information to the recommender system about what he or she likes and dislikes. This is a fundamental task that is strictly necessary to provide personalized recommendations. If the system has no specific knowledge about the active user, then it can only provide the same recommendations that would be delivered to an "average" user.
- *Express self*: Some users may not care about the recommendations at all. Rather, what is important to them is that they be allowed to contribute with their ratings and express their opinions and beliefs. The user satisfaction for that activity can still act as leverage, resulting in the user's continued loyalty to the application (as we mentioned prior, in discussing the service provider's motivations).
- *Help others*: Some users are happy to contribute with information, e.g., their evaluation of items (ratings), because they believe that the community benefits from their contribution. This could be a major motivation for entering information into a recommender system that is not used routinely. For instance, with an automobile RS, a user who has already purchased a new car is aware that the rating entered in the system is more likely to be useful to other users rather than to oneself, the next time a new-car-purchase is contemplated.
- *Influence others*: In Web-based RSs, there are users whose main goal is to explicitly influence other users into purchasing particular products. As a matter of fact, there are also malicious users that may use the system simply to promote or penalize certain items (see Chap. 28).

As these various points indicate, the role of an RS within an information system can be quite diverse. This diversity calls for the exploitation of a range of different knowledge sources and techniques. In the next two sections, we discuss the data that an RS manages and the core technique used to identify the right recommendations.

1.3 Data and Knowledge Sources

RSs are information processing systems that actively gather various kinds of data in order to build their recommendations. Data is primarily about the items to suggest and the users who will receive these recommendations. But, since the data and knowledge sources available for recommender systems can be very diverse, ultimately, whether it can be exploited or not depends on the recommendation technique (see also Sect. 1.4). This will become clearer in the various chapters included in this handbook.

In general, there are recommendation techniques that are knowledge-poor, namely, that use very simple and basic data, such as user ratings or evaluations for items (Chaps. 2 and 3). Other techniques are much more knowledge-dependent, in that they use ontological descriptions of the users or the items (Chap. 4), constraints (Chap. 5), or social relations and activities of the users (Chaps. 15 and 17). In any case, as a general classification, data used by RSs refers to three kinds of objects: items, users, and transactions, that is, relations between the users and the items.

Items Items are the objects that are recommended. Items may be characterized by their complexity and their value or utility. The value of an item may be positive if the item is useful to the user, or negative if the item is not appropriate and the user made the wrong decision when selecting it. We note that when a user is acquiring an item, one will always incur in a cost which includes the cognitive cost of searching for the item and the real monetary cost eventually paid for the item.

For instance, the designer of a news RS must take into account the complexity of a news item, i.e., its structure, the textual representation, and the time-dependent importance of any news item. But at the same time, the RS designer must understand that even if the user is not paying for reading news, there is always a cognitive cost associated with searching and reading news items. If a selected item is relevant to the user, this cost is dominated by the benefit of having acquired useful information. Whereas if the item is not relevant, the net value of that item for the user, and its recommendation, is negative. In other domains, e.g., cars, or financial investments, the true monetary cost of the items becomes an important element to consider when selecting the most appropriate recommendation approach.

Items with low complexity and value are: news, webpages, books, CDs, and movies. Items with larger complexity and value are: digital cameras, mobile phones, PCs, etc. The most complex items that have been considered are insurance policies, financial investments, travel, and jobs [39].

RSs, according to their core technology, can use a range of properties and features of the items. For example in a movie recommender system, the genre (comedy, thriller, etc.), as well as the director and actors, can be used to describe a movie and to learn how the utility of an item depends on its features. Items can be represented using various information and representation approaches, e.g., in a minimalist way as a single ID code, or in a richer form, as a set of attributes, and even as a concept in an ontological representation of the domain (Chap. 4).

Users Users of an RS, as mentioned above, may have very diverse goals and characteristics. In order to personalize the recommendations and the human-computer interaction, RSs exploit a range of information about the users. This information can be structured in various ways, and again, the selection of what information to model depends on the recommendation technique.

For instance, in collaborative filtering, users are modeled as a simple list containing the ratings provided by the user for certain items. In a demographic RS, sociodemographic attributes such as age, gender, profession, and education, are used. User data is said to constitute the user model [12, 22]. The user model profiles the user, i.e., encodes her preferences and needs. Various user modeling approaches have been used and, in a certain sense, an RS can be viewed as a tool that generates recommendations by building and exploiting user models [10, 11]. Since no personalization is possible without a convenient user model the user model will always play a central role. For instance, in reconsidering a collaborative filtering approach, the user is either profiled directly by its ratings of items or, using these ratings, the system derives a vector of factor values where users differ in how each factor weights in their model (Chaps. 2 and 3).

Users can also be described by their behavior pattern data, for example, site browsing patterns (in a Web-based recommender system) [54], or travel search patterns (in a travel recommender system) [35]. Moreover, user data may include relations between users such as the trust level of these relations between users (Chap. 16). An RS might utilize this information to recommend items to users that were preferred by similar or trusted users.

Transactions We generically refer to a transaction as a recorded interaction between a user and the RS. Transactions are log-like data that store important information generated during the human-computer interaction and which are useful for the recommendation generation algorithm that the system is using. For instance, a transaction log may contain a reference to the item selected by the user and a description of the context (e.g., the user goal/query) for that particular recommendation. If available, that transaction may also include explicit feedback that the user has provided, such as the rating for the selected item.

In fact, ratings are the most popular form of transaction data that an RS collects. These ratings may be collected explicitly or implicitly. In the explicit collection of ratings, the user is asked to provide an opinion about an item on a rating scale. According to [47], ratings can take on a variety of forms:

- Numerical ratings such as the 1–5 stars provided in the book recommender associated with Amazon.com.
- Ordinal ratings, such as “strongly agree, agree, neutral, disagree, strongly disagree” where the user is asked to select the term that best indicates his or her opinion regarding an item (usually via questionnaire).
- Binary ratings that model choices in which the user is simply asked to decide if a certain item is good or bad.

- Unary ratings can indicate that a user has observed or purchased an item, or otherwise rated the item positively. In such cases, the absence of a rating indicates that we have no information relating the user to the item (perhaps the user purchased the item elsewhere).

Another form of user evaluation consists of tags associated by the user with the items that the system presents. For instance, on MovieLens (<http://movielens.umn.edu>), RS tags represent how MovieLens users feel about a movie, e.g.: “too long,” or “acting.”

In transactions that collect implicit ratings, the system aims to infer the user’s opinion based on the user’s actions. For example, if a user enters the keyword “Yoga” at Amazon.com, a long list of books will be provided. In return, the user may click on a certain book on the list in order to receive additional information. At this point, the system may infer that the user is somewhat interested in that book.

In conversational systems, i.e., systems that support an interactive process, the transaction model is more refined. In these systems, user requests alternate with system actions (see Chaps. 10 and 18). That is, the user may request a recommendation and the system may produce a suggestion list. But it can also request additional user preferences to provide the user with better, more refined results. Here, in the transaction model, the system collects the various requests-responses, and may eventually learn to modify its interaction strategy by observing the outcome of the recommendation process [35].

1.4 Recommendation Techniques

In order to implement its core function, identifying useful items for the user, a RS must *predict* that an item is worth recommending. In order to do this, the system must be able to predict the utility of some items, or at least compare the utility of some items, and then decide which items to recommend based on this comparison. The prediction step may not be explicit in the recommendation algorithm but we can still apply this unifying model to describe the general role of an RS. Here, our goal is to provide the reader with a unifying perspective rather than an account of all the different recommendation approaches that will be illustrated in this handbook.

To illustrate the prediction step of an RS, consider for instance, a simple and non-personalized recommendation algorithm that recommends only the most popular songs. The rationale for using this approach is that in the absence of more precise information about the user’s preferences, a popular song, i.e., one that is liked (high utility) by many users, will also most-likely appeal to a generic user, or at least with a higher likelihood than another randomly selected song. Hence, the utility of such popular songs is predicted to be reasonably high for this generic user.

This view of the core recommendation computation as the prediction of the utility of an item for a user has been suggested in [2] and recently updated in [44]. Both papers model this degree of utility of the user u for the item i as a (real valued)

function $R(u, i)$, as is normally done in collaborative filtering by considering the ratings of users for items. Then, the fundamental task of a collaborative filtering RS is to predict the value of R over pairs of users and items, or in other words, to compute $\hat{R}(u, i)$, where we denote with \hat{R} the estimation, computed by the RS, of the true function R . Consequently, having computed this prediction for the active user u on a set of items, i.e., $\hat{R}(u, i_1), \dots, \hat{R}(u, i_N)$, the system will recommend the items i_{j_1}, \dots, i_{j_K} ($K \leq N$) with the largest predicted utility. K is typically a small number, that is, much smaller than the cardinality of the item data set or the items on which a user utility prediction can be computed, i.e., RSs “filter” the items that are recommended to users.

As mentioned above, some recommender systems do not fully estimate the utility before making a recommendation, but they may apply some heuristics to hypothesize that an item may be of use to a user. This is typical, for instance, in knowledge-based systems. These utility predictions are computed with specific algorithms (see below) and use various kinds of knowledge about users, items, and the utility function itself (see Sect. 1.3) [17]. For instance, the system may assume that the utility function is Boolean and therefore it will just determine whether an item is or is not useful for the user. Consequently, assuming that there is some available knowledge, or possibly none, about the user who is requesting the recommendation, as well as knowledge about items, and other users who received recommendations, the system will leverage this knowledge with an appropriate algorithm to generate various utility predictions and hence recommendations [17].

It is also important to note that sometimes the user utility for an item is observed to depend on other variables, which we generically call “contextual” [44]. For instance, the utility of an item for a user can be influenced by the domain knowledge of the user (e.g., expert versus beginning users of a digital camera), or can depend on the time when the recommendation is requested. Equally, users may be more interested in items (e.g., restaurant) closer to their current location. Consequently, the recommendations must be adapted to these specific additional details and as a result it becomes increasingly more difficult to correctly estimate what the right recommendations are.

This handbook presents several different types of recommender systems that vary in terms of the addressed domain and the knowledge used, but especially with regard to the recommendation algorithm, i.e., how the prediction of the utility of a recommendation is made, as was mentioned at the beginning of this section. Other differences relate to how the recommendations are finally assembled and presented to the user in response to user requests. These aspects are discussed as well, later in this introduction.

To provide an initial overview of the different types of RSs, we want to quote a taxonomy provided by Burke [17] that has become a classical way of distinguishing between recommender systems and referring to them. Burke [17] distinguishes between six different classes of recommendation approaches:

Content-Based The system learns to recommend items that are similar to the ones that the user liked in the past. The similarity of items is calculated based on the

features associated with the compared items. For example, if a user has positively rated a movie that belongs to the comedy genre, then the system can learn to recommend other movies from this genre [33].

Classic content-based recommendation techniques aim at matching the attributes of the user profile against the attributes of the items. In most cases, the items' attributes are simply keywords that are extracted from the items' descriptions. Semantic indexing techniques represent the item and user profiles using concepts instead of keywords. Chapter 4 presents a comprehensive survey of semantic indexing techniques to overcome the main problems of classical keyword-based systems. The authors presents two main groups of semantic indexing techniques: top-down and bottom-up. Techniques in the former group rely on the integration of external knowledge sources, such as: ontologies, encyclopedic knowledge (such as Wikipedia) and data from the Linked Data cloud, while techniques in the latter group rely on a lightweight semantic representation based on the hypothesis that the meaning of words depends on their usage in large corpora of textual documents. Chapter 4 demonstrates how to utilize semantic approaches to realize a new generation of semantic content-based recommender systems, by providing a description of their main potentials and limitations.

Collaborative Filtering The original and most simple implementation of this approach [24] makes recommendations to the active user based on items that other users with similar tastes liked in the past. The similarity in taste of two users is calculated based on the similarity in the rating history of the users. This is the reason why [48] refers to collaborative filtering as “people-to-people correlation.” Collaborative filtering is considered to be the most popular and widely implemented technique in RS.

Chapter 2 presents a comprehensive survey of neighborhood-based methods for collaborative filtering. Neighborhood-based methods focus on relationships between items or, alternatively, between users. An item-item approach models the preference of a user to an item based on ratings of similar items by the same user. Neighborhood-based methods benefit from considerable popularity due to their simplicity, efficiency, and ability to produce accurate and personalized recommendations. Chapter 2 describes the main benefits of such methods, as well as their principal characteristics.

Moreover, Chap.2 addresses the essential decisions that are required while implementing a neighborhood-based recommender system, and gives practical information on how to make such decisions. Perhaps the decision that has the greatest impact on the rating prediction and computational performance of the recommender system is the choice between a user-based and an item-based method. In typical commercial recommender systems where the number of users exceeds the number of available items, item-based approaches should be preferred since they provide more accurate recommendations, while being more computationally efficient and requiring less frequent updates. On the other hand, user-based methods usually provide more original recommendations, which may lead users to a more satisfying experience [21].

Finally, the problems of sparsity and limited coverage, often observed in large commercial recommender systems, are discussed by exploring two research directions: dimensionality reduction and graph-based techniques. Dimensionality reduction provides a compact representation of users and items that captures their most significant features. An advantage of such an approach is that it allows for obtaining meaningful relations between pairs of users or items, even though these users have rated different items, or these items were rated by different users. On the other hand, graph-based techniques exploit the transitive relations in the data. These techniques also avoid the problems of sparsity and limited coverage by evaluating the relationship between users or items that are not directly connected. However, unlike dimensionality reduction, graph-based methods also preserve some of the “local” relations in the data.

Chapter 3 presents several recent extensions available for building CF recommenders. Specifically, the authors discuss latent factor models, such as matrix factorization (e.g., Singular Value Decomposition (SVD)). These methods transform both items and users to the same latent factor space. The latent space is then used to explain ratings by characterizing both products and users in term of factors automatically inferred from user feedback. The authors elucidate how SVD can handle additional features of the data, including implicit feedback and temporal information. They also describe techniques to address shortcomings of neighborhood techniques by suggesting more rigorous formulations using global optimization techniques. Utilizing such techniques makes it possible to lift the limit on neighborhood size and to address implicit feedback and temporal dynamics. The resulting accuracy is close to that of matrix factorization models, while offering a number of practical advantages.

Demographic This type of system recommends items based on the demographic profile of the user [13]. The assumption is that different recommendations should be generated for different demographic niches. Many websites adopt simple and effective personalization solutions based on demographics. For example, users are dispatched to particular websites based on their language or country. Or, suggestions may be customized according to the age of the user. While these approaches have been quite popular in the marketing literature, there has been relatively little proper RS research on demographic systems.

Knowledge-Based Knowledge-based systems recommend items based on specific domain knowledge about how certain item features meet users’ needs and preferences and, ultimately, how the item is useful for the user. Notable knowledge-based recommender systems are case-based [15, 19, 45]. In these systems, a similarity function estimates how much the user’s needs (problem description) match the recommendations (solutions of the problem). Here, the similarity score can be directly interpreted as the utility of the recommendation for the user. Knowledge-based systems tend to work better than others at the beginning of their deployment but if they are not equipped with learning components, they may be surpassed by other shallow methods that can exploit the logs of the human/computer interaction (as in CF).

Constraint-based systems are another type of knowledge-based RS (Chap. 5). In terms of used knowledge, both systems are similar: user requirements are collected, repairs for inconsistent requirements are automatically proposed in situations where no solutions could be found, and recommendation results are explained. The major difference lies in the way solutions are calculated. Case-based recommenders determine recommendations on the basis of similarity metrics whereas constraint-based recommenders predominantly exploit predefined knowledge bases that contain explicit rules about how to relate customer requirements with item features.

Chapter 5 reviews constraint-based recommendation approaches and provide an overview of technologies for the development of knowledge bases for constraint-based recommenders since appropriate tool support can be crucial in practical settings. The authors show that constraint-based methods are particularly well suited for recommending complex products such as financial services or electronic consumer goods.

Moreover, Chap. 5 presents possible forms of user interaction that are supported by constraint-based recommender applications, report scenarios in which constraint-based recommenders have been successfully applied, and review different technical solution approaches.

Community-Based This type of system recommends items based on the preferences of the user's friends. This technique follows the epigram, "Tell me who your friends are, and I will tell you who you are" [4, 9]. Evidence suggests that people tend to rely more on recommendations from their friends than on recommendations from similar but anonymous individuals [52]. This observation, combined with the growing popularity of open social networks, is generating a rising interest in community-based systems or, as they are usually referred, social recommender systems [23]. This type of RS models and acquires information about the social relations of the users and the preferences of the user's friends. The recommendation is based on ratings that were provided by the user's friends. In fact these RSs are following the rise of social-networks and enable a simple and comprehensive acquisition of data related to the social relations of the users.

Hybrid Recommender Systems These RSs are based on the combination of the above mentioned techniques. A hybrid system combining techniques A and B tries to use the advantages of A to fix the disadvantages of B. For instance, CF methods suffer from new-item problems, or, that they cannot recommend items that have no ratings. This does not limit content-based approaches since the prediction for new items is based on their description (features) that are typically easily available. Given two (or more) basic RSs techniques, several ways have been proposed for combining them to create a new hybrid system (see [17] for the precise descriptions).

As we have already mentioned, the context of the user when he or she is seeking a recommendation can be used to better personalize the output of the system. For example, in a temporal context, vacation recommendations in winter should be very different from those provided in summer [8]. Or a restaurant recommendation for a Saturday evening with one's friends should be different from that suggested for a workday lunch with co-workers.

Chapter 6 reviews the topic of context-aware recommender systems (CARS). It presents the general notion of context and how it can be modeled in RSs. As it discusses the possibilities of combining several context-aware recommendation techniques into a single unified approach, the authors also provide a case study of one such combined approach.

Three popular different algorithmic paradigms for incorporating contextual information into the recommendation process are discussed: reduction-based (prefiltering), contextual post filtering, and context modeling. In reduction-based (prefiltering) methods, only the information that matches the current usage context, e.g., the ratings for items evaluated in the same context, are used to compute the recommendations. In contextual post filtering, the recommendation algorithm ignores the context information. The output of the algorithm is filtered/adjusted to include only the recommendations that are relevant in the target context. In the contextual modeling, the more sophisticated of the three approaches, context data is explicitly used in the prediction model.

Recommendation tasks can be solved with the help of techniques that were developed in the field of Data Mining. Chapter 7, presents an overview of the main Data Mining techniques used in the context of Recommender Systems and presents cases where these techniques have been successfully applied. In particular, it discusses the following techniques: preprocessing techniques such as sampling or dimensionality reduction; classification techniques, such as Bayesian Networks, Decision Trees and Support Vector Machines; clustering techniques such as k -means; and finally association rules.

1.5 Recommender Systems Evaluation

Recommender systems research is being conducted with a strong emphasis on practice and commercial applications. One very important issue related to the practical side of RS deployment is the necessity of evaluating the quality and value of the systems. Evaluation is required at different stages of the system's life cycle and for various purposes [1, 26]. At design time, evaluation is required to verify the selection of the appropriate recommender approach. In the design phase, evaluation should be implemented off-line and the recommendation algorithms, i.e., their computed recommendations, are compared with the stored user interactions. An off-line evaluation consists of running several algorithms on the same datasets of user interactions (e.g., ratings) and comparing their performances. This type of evaluation is usually conducted on existing public benchmark data if appropriate data is available, or, otherwise, on collected data. The design of the off-line experiments should follow known experiment design practices [6] in order to ensure reliable results. Off-line experiments can measure the quality of the chosen algorithm in fulfilling its recommendation task. However, such evaluation cannot provide any insight about the user satisfaction, acceptance or experience with the system. The algorithms might be very accurate in solving the core recommendation

problem, i.e., predicting user ratings, but for some other reason the system may not be accepted by users, for example, because the performance of the system was not as expected.

Therefore, a user-centric evaluation is also required. It can be performed online after the system has been launched, or as a focused user study. During on-line evaluation, real users interact with the system without being aware of the full nature of the experiment running in the background. It is possible to run various versions of the algorithms on different groups of users for comparison and analysis of the system logs in order to enhance system performance. In addition, most of the algorithms include parameters, such as weight thresholds, the number of neighbors, etc., requiring constant adjustment and calibration.

Focused user studies are conducted when the on-line evaluation is not feasible or too risky. In this type of evaluation, a controlled experiment is planned where a small group of users are asked to perform different tasks with various versions of the system. It is then possible to analyze the user's performance and to distribute questionnaires so that users may report on their experience. In such experiments, it is possible to collect both quantitative and qualitative information about the systems.

In recent years there has been an increased interest in user-centric evaluation procedures and metric for recommender systems. Researchers realized that recommender systems' goals extend beyond the accuracy of the algorithms [30] as tools to provide a helpful and enjoyable, personalized experience that leads to user retention and satisfaction. This approach broadened the range of evaluated aspects of an RS to include aspects such as the form of preference elicitation, the presentation of the recommended results (e.g., one top item, top N items, or predicted ratings), and finally, the evaluation of the explanations provided to the users. Explanations may serve a few goals: the most popular is the justification of results, i.e., explaining to the user why the system decided to recommend a specific item. Other goals may include increasing trust in the system, persuading the user to purchase the recommended item, and helping a user with their decision making. When designing the evaluation of the recommendation explanation, it is important to identify the goal of the explanation and adjust a suitable metric for measuring it.

Chapter 8 details the three previously mentioned types of experiments that can be conducted in order to evaluate recommender systems, namely, off-line, on-line and user studies. It presents their advantages and disadvantages, and defines guidelines for choosing the methods for evaluating them by considering the properties that are to be evaluated. Unlike existing discussions of evaluation in the literature that usually focuses mainly on the accuracy of an algorithm's prediction [26] and related measures, this chapter is unique in its approach to the evaluation discussion since it focuses on property-directed evaluation. It provides a large set of properties (other than accuracy) that are relevant to the system's success. For each of the properties, the appropriate type of experiment and relevant measures are suggested. Among the list of properties are: coverage, cold start, confidence, trust, novelty, risk, and serendipity. The chapter describes the difficulties and pitfalls of each of the properties and guidelines for the selection of the suitable evaluation type and properties for a given recommendation task ad system.

Chapter 9 highlights the importance of user-centric evaluation. It provides detailed and practical guidelines of how to conduct user-centric experiments in order to evaluate the user's experience with the system. The chapter first presents a theoretical user-centric evaluation framework [29] that maps aspects of recommender systems and their interaction with the users that should be evaluated. Then, it provides practical guidelines for students and researchers for conducting user experiments. It presents tips for stating hypotheses, recruiting participants, design of the experiments and statistical analysis of the results. Finally, the chapter provides many examples of actual systems' evaluations from the relevant literature.

Chapter 10 tackles an additional aspect of the user-centric approach to evaluation and highlights an important aspect of RS: explanations of the recommendation results to the users. It examines the reasons that make an evaluation "good" and the effects that explanations might have on RS acceptance. The chapter first explains the interaction between the recommender system and the explanation in terms of preference elicitation methods and the presentation of results, as well as the recommendation algorithm. Then, explanation styles are described, along with examples of explanation in existing systems. The goals of explanations are listed from which metrics that measure the success of explanations in achieving these goals are described. The chapter concludes with challenges related to explanations. This includes the context in which explanations should be shown, and a major challenge in evaluating the interaction between acceptance of recommendation and explanations, as well as how to assure that explanations are indeed helpful and do not lead users to make poor decisions.

1.6 Recommender Systems Applications

Recommender systems research, aside from its theoretical contribution, is generally aimed at practically improving industrial RSs and involves research about various practical aspects that apply to the implementation of the systems. Indeed, an RS is an example of large scale usage of machine learning and data mining algorithms in commercial practice [3]. The common interest in the field, both from the research community and from the industry has leveraged the availability of data for research on one hand, and the evolvement of enhanced algorithms on the other hand. Practical related research in RSs examines aspects that are relevant to different stages in the life cycle of an RS, namely, the design of the system, its implementation, evaluation, maintenance and enhancement during system operation. The Netflix Prize announced in 2006, described in Chap. 11, was an important event for the recommender systems research community and industry, and their mutual interaction. It highlighted the importance of the recommendation of items to users and accelerated the development of many new data mining recommendation techniques. Even though the Netflix Prize initiated a lot of research activities, the prize was a simplification of the full recommendation problem. It consisted of predicting user's ratings while optimizing the Root Mean Square Error (RMSE)

between the predicted and actual ratings. Chapter 11 describes lessons learned from the prize awarded in 2009 and provides insights about industrial setting of RSs using the Netflix system as a case study for real-world recommender systems. In addition, the chapter provides the industry perspective about RSs implementation issues that deserve attention while developing a real-world RS. It describes the data centric approach used at Netflix for selecting the best model for a problem in order to provide an optimized personalized experience to the user. In addition, the chapter highlights the need for a suitable scalable system architecture that can support the development and evaluation process of innovative algorithms and deliver recommendations based on large volumes of available data.

The first factor to consider while designing an RS is the application's domain as it has a major effect on the algorithmic approach that should be taken. Montaner et al. [39] provide a taxonomy of RSs and classify existing RS applications to specific application domains. Based on these specific application domains, we define more general classes of domains for the most common recommender systems applications:

- Entertainment—recommendations for movies, music, games, and IPTV.
- Content—personalized newspapers, recommendation for documents, recommendations of webpages, e-learning applications, and e-mail filters.
- E-commerce—recommendations of products to buy such as books, cameras, PCs etc. for consumers.
- Services—recommendations of travel services, recommendation of experts for consultation, recommendation of houses to rent, or matchmaking services.
- Social—recommendation of people in social networks, and recommendations of content social media content such as tweets, Facebook feeds, LinkedIn updates, and others.

As recommender systems become more popular, interest is roused in the potential advantages of new and diverse applications, such as recommending insurance riders, or recommending questions for question-answering systems. As the above list cannot cover all the application domains that are now being addressed by RS techniques: it gives only an initial description of the various types of application domains.

The developer of an RS for a certain application domain should understand the specific facets of the domain, its requirements, application challenges and limitations. Only after analyzing these factors can one be able to select the optimal recommender algorithm and to design an effective human-computer interaction. In the current version of the handbook, some of the chapters in this section describe applications of recommender systems in specific domains. Each of these chapters describes the requirements of an RS for a specific domain, its precise challenges and the suitable technologies and algorithms for addressing them.

One detailed example of an RS designed for a specific domain is described in Chap. 12, which deals with recommender systems for technology-enhanced learning (TEL). TEL, which generally covers technologies that support all forms of teaching and learning activities, aims at designing, developing and testing new methods and

technologies to enhance learning practices of both individuals and organizations. As the digital way of education and learning is becoming more popular, the need to integrate the personalization of content into the learning process and the available data for assessing the quality of the algorithms has created opportunities for the rise of the popularity of TEL RSs. Since TEL may benefit greatly from integrating recommender systems technology to personalize the learning process and adjust it to the user's former knowledge, abilities and preferences [55], there is a significant increase in RSs applied to TEL. The chapter presents an extensive survey of RSs for TEL covering 82 systems from 35 countries categorizing them using a classification framework consisting of three main categories, namely: Supported Tasks, Approach, and Operation. An analysis of the 82 systems using the framework resulted in seven clusters of TEL RSs, where each cluster represents a unique form of contribution to the field. The chapter aims at providing an overview about the TEL RS field in order to standardize evaluation settings and measures, as well as providing guidelines for the application of RS technology to TEL.

Yet another popular domain for recommendation is the recommendation of music, presented in Chap. 13. Unique features of music items that pose various challenges for recommendations should be considered when designing and evaluating RSs for music. Such challenges include, for example, the short time that it takes a user to gain an opinion about a recommended item, as compared to a movie or a book, or the fact that the same item can be recommended many times. In addition, music can be recommended as a single item, a playlist, and abstracted by genre, performer, or band. Music RSs, as opposed to many other domains, rely heavily on content-based recommendation which implies specific challenges to the domain [37].

Some new application domains for recommendation evolved with the emergence of new technologies that became very popular. One example, detailed in Chap. 14, is the evolution of mobile technology that accelerated the development of specific RSs for mobile devices utilizing the special capabilities of the devices (e.g., the GPS). Chapter 14 reviews the main components of a location based mobile recommender system. While there are various application domains of mobile context aware RSs (as described in Chap. 6) that may benefit from mobile sensors in order to enrich their users' profiles, Chap. 14 highlights mobile location-based recommendations. Such RS applications recommend places and venues based on the user's location and history of behaviors and preferences. The chapter describes the algorithms that have been applied to recommending venues, and the evaluation procedures for assessing the quality of the recommendations. Looking into the future, the authors suggest additional location based applications such as recommendations for cab drivers on pick-up locations, or recommending where to locate a retail store.

Another example of new RSs that emerged with new technologies are recommender systems related to the social web, and specifically those that target the social media domain. With the rise of social networks (e.g., Facebook, LinkedIn, Tweeter, Flickr, and others), users are overloaded with information, activities and interactions. Social recommender systems are RSs that aim at assisting the user in identifying relevant content (e.g., tweets, feeds or images), and engage only in

relevant activities and interactions (e.g., discussions, or comments). Apart from the RSs that have been developed to be dedicated to social media, recommender systems of other domains can benefit from the new types of data that social media introduces about users to enhance the quality of standard RSs [23]. The term Social RS covers many types of RSs that are relevant to the social media platform in which Chap. 15 describes two main types: recommendations of social media content and recommendations of people. For recommendations of social content, the chapter reviews various social content media domains, and provides a detailed case study and insights learned from a recommender system operated in the enterprise which suggests mixed social media items. The chapter lists three different types of people recommendations, namely, the recommendation of familiar people (e.g., classmates, family members) that are not connected in the network; recommendations of interesting people (to connect with, or follow), and recommendations of strangers (to date, to hire, or for various other purposes). It explains the complexity of people-recommendation and lists key topics that should be considered and should be further investigated. The list includes: the need for explanation, privacy concerns, social relationships, trust and reputation, as well as the need to define special evaluation measures.

Chapter 16 highlights a special form of social recommendations, the reciprocal people-to-people recommendations that require the two parties involved in a recommendation to be satisfied. Some examples are: dating recommendations, Human Resource recommendations (recommendations of employees and employers), and recommending groupings of students for learning groups. Besides the unique reciprocal manner of satisfaction that is required from both parties, the chapter highlights other differences between traditional and reciprocal recommendations, including the users' willingness to provide explicit feedback, the fact that users are usually engaged with the system for a long term, and the requirement not to overload users with recommendations.

Within the chapter, the authors provide an overview regarding existing reciprocal recommender systems and demonstrate how the special requirements of reciprocity are considered in system design. A detailed case study, specifically in online dating recommendation, is presented and includes the recommendation hybrid content-collaborative algorithm and its evaluation. One interesting insight shown is that implicit feedback is more effective than explicit for these types of systems.

The social Web is also exploited by modern search engines that rely on recommendation techniques to address Web search challenges and to implement advanced search features. Specifically, various engines attempt to apply some form of personalization and collaboration by generating results to a user query that are not only relevant to the query terms but are also tailored to the user's search history, reputation, and preferences as inferred from the user's own previous activities on the social Web.

Chapter 17 discusses the research goals of Information Retrieval (IR) and personalized Web search from the RS perspective. The authors illustrate how techniques that originated in recent RS research may be applied to address search engine challenges. This chapter focuses on two promising ideas for search engine

improvement: personalization and collaboration. It describes a number of different approaches to personalizing Web searches by exploiting user preferences and context information to affect search results. In addition, the chapter discusses recent work in the area of collaborative information retrieval, which attempts to take advantage of the potential for cooperation between friends, colleagues or users with similar needs in implementing a variety of information-seeking tasks. This new line of research, termed social search, benefits from the social medium property of the Web in providing search results that are affected by the experience and preferences of similar users. The authors foresee a convergence of recommender systems and search engines, where a search engine will provide a unique platform for modern recommendation technologies. The authors believe that integrating these sources in search engine algorithms, along with a proactive manner of search experience that strives to understand the users' needs, would result in highly satisfied users that are able to receive the right information at the right time. Another trend that is affecting search engines is the rise of search activities through mobile devices. This introduces new constraints to search and discovery interfaces, but also brings about opportunities for innovations using the mobile sensors that allow enhanced personalization.

1.7 Recommender Systems and Human Computer Interaction

As we have illustrated in the previous sections, researchers have been chiefly concerned with designing a range of technical solutions and leveraging various sources of knowledge to achieve better predictions about what is liked and how much it is liked by the target user. The underlying assumption behind this research activity is that simply presenting these correct recommendations, or the best options, is not sufficient. In other words, the recommendations should speak for themselves, and the user should definitely accept the recommendations if they are correct. This is clearly an overly simplified account of the recommendation problem and the delivery of recommendations is not so straightforward.

In practice, users need recommendations because they do not have enough knowledge to make an autonomous decision. Consequently, it may not be easy for them to evaluate the proposed recommendation. Hence, various researchers have tried to understand the factors that lead to the acceptance of a recommendation by a given user [5, 20, 25, 38, 50, 53].

Swearingen and Sinha [53] were among the first to point out that the effectiveness of an RS is dependent on factors that go beyond the quality of the prediction algorithm. In fact, the recommender must also convince users to try (or read, buy, listen, watch, etc.) the recommended items. This, of course, depends on the individual characteristics of the selected items, and therefore on the recommendation algorithm. The process also depends, however, on the particular human/computer

interaction supported by the system when the items are presented, compared, and explained. Swearingen and Sinha [53] found that from a user's perspective, an effective recommender system must inspire trust in the system and it must have a system logic that is at least somewhat transparent. Additionally, the authors note that it should point users towards new, not-yet-experienced items, and should provide details about recommended items, including pictures and community ratings, and finally, it should present ways to refine recommendations.

Swearingen and Sinha [53] and other similarly oriented researchers do not diminish the importance of the recommendation algorithm, but claim that its effectiveness should not be evaluated only in terms of the accuracy of the prediction, i.e., with standard and popular IR metrics, such as Mean Absolute Error (MAE), precision, or Normalized Discounted Cumulative Gain (NDCG) (see Chap. 8). Other dimensions should be measured that relate to the acceptance of the recommender system and its recommendations. These ideas have been remarkably well presented and discussed also by McNee et al. [38]. In that work, the authors propose user-centric directions for evaluating recommender systems, including: the similarity of recommendation lists, recommendation serendipity, and the importance of user needs and expectations in a recommender.

Recommender systems collect tremendous amounts of user data that are necessary for the recommendation purposes. However, the availability of this data may result in this data being used in a way that violates the end-user's expectations of privacy, especially if it is accessed by untrusted parties or misused by malicious agents. Chapter 19 presents the latest in privacy enhanced recommendations. The authors analyze the risks to user privacy imposed by recommender systems, survey the existing solutions, and discuss the privacy implications for the users of the recommenders. In particular, the authors describe several architectures that preserve user privacy by using various decentralized solutions that eliminate a single repository of user modeling data, which would otherwise be the target for malicious attacks on the recommender. In addition, the authors present algorithmic solutions, which either perturb the original user modeling data or apply formal encryption methods. These assure that, even if accessed by an untrusted party, only modified or encrypted user data would be exposed rather than the original data. Lastly, the policy driven solutions are described. These solutions address directives and legislation initiatives that limit the storage, transfer, and exploitation of personal user data.

An essential goal of recommender systems is to help users make better choices [18]. Thus, it is important to understand how people make choices and how the human decision making process can be supported. Chapter 18 begins with a compact overview of the psychology of everyday choice and decision making that is based on a large literature of psychological research and formulated so as to be relevant and accessible to recommender systems research community. The authors explain how recommender systems can be viewed as one of many available tools for facilitating choice. Then, the authors provide a high-level overview of strategies for helping people make better choices, indicating how recommender systems fit into the greater picture of choice. In addition, the main functionalities of RSs are presented: eliciting information to construct preference models, narrowing down a large set of options,

helping users choose among a small set of recommended options, and helping users to explore large spaces of options. The authors show how an understanding of human decision making can illuminate research and practice concerning these processes.

As discussed in previous sections, recommender systems often utilize sophisticated algorithms to make recommendations. However, the RS cannot assume the advice provided by a system will always be accepted by its users. Whether a recommendation is seen as credible advice and actually taken into account not only depends on users' perceptions of the recommendation but also of the system as an advice-giver.

In Chap. 20 the authors stress that a recommendation is seen as credible advice and is actually taken into account not only because of the user's perceptions of the recommendation but also due to the fundamental role of the system which is perceived as an advice-giver. Indeed, the literature about persuasion suggests that people are likely to accept recommendations from credible sources and we therefore conclude that the credibility of the RS is vital to increasing the likelihood of recommendation acceptance. Hence, the authors discuss how the credibility of RSs can be enhanced, providing a synopsis of credibility-related research.

Chapter 20 reviews the existing literature on source factors in the context of human-human, human-technology, and human-recommender system interactions. It also discusses system credibility evaluation in light of the increasing popularity of social technology. Source characteristics which have been studied in the context of human and technology interaction, and particularly, in the recommender systems realm are discussed as well. Finally, Chap. 20 concludes that many social cues that have been identified as influential in other contexts have yet to be implemented and tested with respect to recommender systems.

Personality accounts for the most important way in which individuals differ in their enduring emotional, interpersonal, experiential, attitudinal and motivational styles. Studies have shown that personality was especially useful at tackling the issues of the cold start problem and diverse recommendations.

Chapter 21 discusses how personality relates to user preferences and how to use personality in recommender systems. The authors present the Five Factor Model (FFM) of personality. This model appears suitable for usage in recommender systems as it can be easily quantified in terms of features corresponding to the main factors. The acquisition of the personality factors for an observed user can be made explicitly through questionnaires or implicitly using machine learning approaches on modalities like social media streams or mobile phone call logs.

1.8 Advanced Topics

It is clear from the previous pages that RS research is evolving in numerous and diverse directions, and new topics are emerging or becoming more important subjects of investigation. The reader is also encouraged to refer to the proceedings

of the last editions of the ACM RecSys conferences and other excellent review papers for additional material [13, 30, 34, 40, 44]. In this handbook, we cover some of these topics. Indeed, some have already been presented, such as: context-aware recommendations (Chap. 6), social-based recommendations (Chap. 15), and reciprocal recommendations (Chap. 16). Other important topics are covered in the last section of this handbook and will now briefly introduce these chapters.

Chapter 22 deals with situations in which the system should recommend information or items that are relevant to a group of users rather than to an individual. For instance, a RS may select television programs for a group to view or a sequence of songs to listen to, based on models of all the group members. Recommending to groups is clearly more complicated than recommending to individuals. Assuming that we know precisely what is good for individual users, the issue is how to combine individual user models. In this chapter, the authors discuss how group recommendation works, what its problems are, and what advances have been made so far.

Chapter 23 discusses the ubiquitous issue of aggregating preferences, criteria or similarities. Normally such aggregation is done by using either the arithmetic mean or maximum/minimum functions. But many other aggregation functions which could deliver flexibility and adaptability, and ultimately more relevant recommendations, are often overlooked. In this chapter the authors review the basics of aggregation functions and their properties and present the most important families, including generalized means, Choquet and Sugeno integrals, ordered weighted averaging, triangular norms and conorms, as well as bipolar aggregation functions. Such functions can model various interactions between the inputs, including conjunctive, disjunctive and mixed behavior.

In Chap. 24, the authors focus on another fundamental problem of RSs: the need to actively look for new data during the operational life of the recommender. This issue is normally neglected on the assumption that there is not much space for controlling the data (e.g., ratings) that the system can collect, since these decisions are taken by the users when visiting the system. Actually, the RS provokes the users with its recommendations and many systems actually explicitly ask for user preferences during the recommendation process. Hence, by tuning the process, users can be pushed to provide a range of different information. Specifically they can be requested to rate particular items since the knowledge of the user's opinions about these items could be estimated as particularly beneficial to the system performance. For instance, the system may be able to provide more diverse recommendations with this additional information, or may improve its prediction accuracy. At this point active learning comes in; it can augment RSs, helping users to become more self-aware of their own likes/dislikes, and lead to more meaningful and useful questions. At the same time, active learning can provide new information to the system that can be analyzed for subsequent recommendations. Hence, applying active learning to RSs enables personalization of the recommending process [36]. This is accomplished by allowing the system to actively influence the items that the user is exposed to (e.g., the items displayed to the user during sign-up or during regular use), as well as by enabling the user to explore his or her interests freely.

Chapter 25 introduces another emerging topic: multi-criteria recommender systems. In the majority of RSs, the utility associated with an item is usually considered a single criterion value, or an overall evaluation or rating of an item by a user. But recently, this assumption has been judged as limited because the suitability of the recommended item for a particular user may depend on several aspects that the user can take into consideration when making his or her choice. The incorporation of multiple criteria that can affect the user's opinions may lead to more effective and accurate recommendations.

Chapter 25 provides an overview of multi-criteria RSs. First, it defines the recommendation problem as a multi-criteria decision-making problem and reviews methods and techniques that can support the implementation of multi-criteria recommenders. Then, it focuses on the category of multi-criteria rating recommender techniques that provide recommendations by modeling the user's utility for an item as a vector of ratings along several criteria. A review of current algorithms that use multi-criteria ratings for calculating the rating prediction and generating recommendations is provided. The chapter concludes with a discussion on open issues and future challenges for these recommenders.

Accurately predicting that the user surely likes a small number of items and suggesting them repeatedly it is not likely to provide a quality service to the user. Likewise, if the suggested items are all quite similar, the system's suggestions will seem repetitive and it may fail to identify peculiar items. Such peculiar items, though they are at the boundary of the user's spectrum of preference, may be unexpected and therefore perceived as more informative or highly desirable to the user.

These topics are discussed in Chap. 26 that is dedicated to novelty and diversity in recommender systems. After having motivated the need for novel and diverse recommendations, the chapter precisely defines these two naturally vague concepts and their relationships. While novelty is a property of the individual recommendations, diversity can only be measured by referring to sets of recommendations, either for a single user or for the full set of users of a recommender system. The quantitative measure of these two properties occupies a substantial amount of the chapter in that this is an important preliminary step prior to focusing on techniques that will improve the system's performance in this respect. This issue is then discussed by presenting a range of techniques aimed at enhancing a recommendation list by editing it with the ultimate goal of increasing the diversification of the recommendations and their novelty. Finally the chapter presents a unifying model that can describe the range of metrics presented thus far. A critical element in this unifying model is that novelty is relative to a context of experience, showing the unexpected relationship between this line of research and that on context-aware recommender systems (see Chap. 6).

As we have referred to already in this introduction, recommender system applications are normally restricted to a particular type of items (movies, CDs, books, etc.). Even though an ecommerce player sells several different types of products (e.g., Amazon), its recommender system treats each product typology independently from the others. This means that if a user has expressed his preferences for a certain category of products (e.g., books), for instance in the form of ratings, this

information is ignored when the system computes recommendations for items in another category (e.g., movies). But, it is natural to expect that preferences in a domain are somewhat correlated to preferences in another domain. So, for instance, if two users have bought very similar sets of books, it is likely that some of the movies bought by one of the users may be good recommendations for the other user.

The idea of relating recommendations in different domains by exploiting user data collected in one domain to produce recommendations in another, is at the base of the research on cross-domain recommender systems. These systems and their underlying techniques are illustrated in Chap.27. It is shown here how leveraging all the user's preferences, collected in several domain specific systems, may be beneficial for generating more comprehensive user models and better recommendations. Cross-domain recommender systems can in fact mitigate the cold-start and sparsity problems in a target domain where not much user data has been collected. Moreover, they can enable cross-selling recommendations, such as building more complex recommendations where items from multiple domains are suggested together (e.g, a movie and a book on a recommended singer). In this chapter, the authors formally define the cross-domain recommendation problem, and try to provide a unifying perspective by merging ideas and approaches which arise in distinct disciplines. The chapter provides an analytical categorization of prior work, and identifies open issues for future research.

The last chapter of this handbook (Chap.28) surveys articles which deal with security issues. This topic has become a major issue in the past few years. The chapter analyzes algorithms designed to generate more robust recommendations, i.e., recommendations that are harder for malicious users to influence. In fact, collaborative recommender systems are dependent on the goodwill of their users, in that there is an implicit assumption that users will interact with the system with the aim of getting good recommendations for themselves while providing useful data for their neighbors. However, users will have a range of purposes in interacting with RSs and in some cases, these purposes may be counter those of the system owner or those of the majority of the user population. Namely, these users may want to damage the website which hosts the recommender, or to influence the recommendations provided to visitors, e.g., to score some items better or worse rather than to arrive at a fair evaluation.

In this chapter, the authors provide a model of efficient attacks, i.e, attacks that can, at a relatively low cost, produce a large impact on system output. Since these attacks may very well be launched against a site, it makes sense to detect them so that countermeasures can be taken as soon as possible. At the same time, researchers have studied a number of algorithms that are intended to robustly withstand attacks and which have lower impact curves relative to efficient attacks. These approaches are also surveyed in this chapter. With the combination of these techniques, researchers have sought not to eliminate attacks, but rather to control their impact to the point in which they are no longer cost-effective.

1.9 Challenges

The list of newly emerging and challenging RS research topics is not limited to those described in the chapters that we have summarized above. Moreover, covering all of them is not within the scope of this introduction. The reader is referred to the final discussion sections that are included in almost all of the chapters published in this handbook for other crucial problems.

Below, we briefly introduce additional challenging topics that we consider important for the development of the research on RSs.

1.9.1 Preference Acquisition and Profiling

A number of open issues are related to the critical stage of acquiring information about the user preferences and generating usable profiles of the users.

It is clear that in many real-world situations, implicit feedback is much more readily available and requires no extra effort on the user's side. For instance, on a web page it is easy to log the users visiting a URL, or clicking on an ad. The system can treat these actions as a form of positive feedback to the displayed items. It makes sense that information about such previous actions contains highly relevant information for predicting future actions. For that reason many recent approaches focus on the use of the more reliable and readily available implicit feedback (see the final discussion in Chap. 11 that refers to this issue). In cases where we have implicit feedback, the recommendation problem becomes the prediction of the probability that a user will interact with a given item. But, standard recommendation formulation in such a setting is not applicable: there is no negative feedback, all the available data is either positive or missing. The missing data includes both items that the user explicitly chose to ignore because they were not appealing and items that would have been perfect recommendations but were never presented to the user. This issue is discussed in Chap. 11 and effective solutions for tackling it must be still developed further.

Notwithstanding the pros of implicit feedback, this data cannot completely substitute the usage of explicitly user-made evaluations, at least for some of the items. In order to make the acquisition of this information more effective, as is illustrated in Chap. 24, a number of techniques for actively collecting preference data from users, especially in the cold start phase, but also in the full life of a system, have been illustrated. It is however still challenging to select, among the various active learning techniques, the one that is more appropriate to a system when dealing with a particular user in a given stage of the system evolution and for a particular goal (e.g., accuracy vs. coverage). More adaptive solutions which can blend, run-time with several elementary approaches, should be investigated. Moreover, their

practical application depends also on their computational cost, which could become prohibitive when multiple strategies must be estimated in real time to find out the best one for that specific scenario.

Hence, in addition to better active learning techniques, simplifying the cognitive cost of preference acquisition is of primary importance. For a recommender system to achieve good recommendation performance, users typically need to provide the system with a certain amount of feedback about their preferences (e.g., in the form of item ratings). This can be an issue in single-rating recommender systems and even more for multi-criteria rating systems that require a more significant level of user involvement, as each user needs to rate an item based on multiple criteria. Therefore, it is important to measure the costs and benefits of adopting alternative rating approaches and scales, and find an optimal solution to meet the needs of both the users and the system designers. For instance, preference disaggregation methods could support the implicit formulation of a preference Multi-Criteria Recommender Systems model (as indicated in Chap. 25).

Another direction of research for easing the preference acquisition process consists of the exploitation of personality, mood and emotions. This is becoming a popular topic, especially because it is clear that more and more techniques will be developed in order to automatically acquire such information. In Chap. 21, the authors stress the challenge of acquiring personality information in a nonintrusive fashion. Nowadays, only the longest questionnaires, which consist of around one hundred questions, can provide an accurate evaluation of the user's personality. Hence, non-intrusive approaches are necessary and the research in this area is just starting. Mining user activity for extracting personality information is an option, but also the fast penetration of portable devices that are life-logging the user's activity can offer a promising platform that is worth exploring.

Another line of research aimed at tackling the cold start problem and reducing the user model elicitation effort is cross-domain recommender systems. These techniques (See Chap. 27) could be used as an alternative path to user preferences' elicitation tools as they are able to build detailed user profiles without the need to collect explicit user assessment of the target domain items. Finally, we want to mention the issue of integrating long-term and short-term user preferences into the process of building a user profile and delivering a recommendation list. Recommender systems may be divided into two classes: those that build a long-term profile, generated by aggregating all the user transaction data collected by the system (e.g., collaborative filtering) and those that are more focused on capturing the ephemeral preferences of the user (e.g., knowledge-based approaches), such as case-based. Obviously both aspects are important and either the precise user task or the availability of items may come into consideration in resolving the preference integration problem. In fact, new research is required to build hybrid models that can correctly decide to what extent to drift toward the contingent user's preferences when there is enough evidence suggesting that the user's short-term preferences are departing from the long-term ones.

1.9.2 Interaction

A major challenge that RS research is now facing is clearly discussed in Chap. 9, in that we still need to broaden the scope of research to the system aspects of a recommender system. This means that aside from the algorithms, which are used to compute the recommendations, the mechanism through which users provide their input and the means by which they receive the systems output, play a significant role and can play an even larger role in determining the success or failure of a recommender system. We still need to better understand the general qualities of alternative solutions to preference elicitation, as we mentioned previously, recommendation presentation and to develop personalized solutions for these phases of the interaction with the system.

It must be observed that while interacting with a recommender system, users make various types of decisions. The most important one is surely selecting an item from the recommendation list. But, before making the final decision, users often have to decide how to explore the information space and what information they must provide to the system. For instance, they could have to select a specific feature (e.g., a camera's size or zoom factor) as search or critiquing criteria, or to select a repair proposal for inconsistent user preferences when interacting with a knowledge-based recommender. Moreover, users often do not know or do not reflect on their preferences beforehand, and the system-supported interaction and visualization contribute to the user construction of their preferences within a specific recommendation scenario. As it has been illustrated in Chap. 18, there are several challenges with the full support of user decision making in a recommender system. Our understanding of the situational context generated by the system and its effect on item selection processes is still incomplete and we need to better connect RS research to psychology and decision making disciplines. While it is clear that RS helps to make decisions, there is still the need for further research that takes theories from decision psychology and cognitive psychology into account when explaining users' preference construction and decision making process in the context of recommender systems.

Considering the user interaction with the recommender system, the topic of explaining the system recommendations still poses a number of interesting and open issues (see Chap. 10). For instance, it is still not completely clear whether explanations bring more overall benefits than risks. In Chap. 10 it is shown that explanations are part of a cyclical process: the explanations affect the acceptance of particular recommendations, the users' mental model of the system, and in turn, this affects the ways users interact with the explanations. But, whether the users are influenced in such a way that their choices are improved is not clear, and explanations may even increase the information overload that the recommendations are supposed to tame. Moreover, while some research has been conducted on explaining recommendations to individual users, explaining recommendations for a group is a much more novel subject. For instance, one might think that accurate predictions of individual satisfaction can also be used to improve the recommender's

transparency: showing how satisfied other group members are could improve users' understanding of the recommendation process and perhaps make it easier to accept items they do not like. However, users need for privacy is likely to conflict with their need for transparency and showing the preferences of other users may move the group discussion on the preferences rather than on the recommended items. We definitely need more research on these topics.

In a discussion about the interaction with recommender systems, we cannot forget the issue of the assessment of the value of the recommendations, which is not only related to what extent the recommended items are liked by the user. For instance, the time value of recommendations, which is partially discussed in the chapter on context-aware recommenders (Chap. 6), refers to the fact that a given set of recommendations may not be applicable forever but there could be a time interval when these items can be recommended. This is clear, for instance, when it comes to news items: people want to be informed about the most recent events and news cannot be meaningfully recommended even 1 day after the initial announcement. The time value of a recommendation is clearly dependent on the novelty and diversity of the recommended items. We still need more theoretical, methodological and algorithmic developments around these aspects. For instance, modeling feature-based novelty in probabilistic terms in order to unify discovery and familiarity models would be an interesting line for future work. Aspects such as the time dimension during which items may recover part of their novelty value, or the variability among users regarding their degree of novelty-seeking are examples of issues that require further research and are mentioned in Chap. 26.

Somewhat connected to the novelty and diversity topic is the issue of achieving an effective tradeoff between exploration and exploitation, which is touched upon in the active learning chapter (Chap. 24). This challenge refers to the fundamental dilemma that a designer must properly tackle, or decide whether to keep recommending items that the system can now identify as good recommendations, given the data currently available to the system, or to further explore user preferences (e.g., asking the user to rate additional, particular items) in order to build newer and possibly better recommendations in the future.

1.9.3 New Recommendation Tasks

The application of recommender systems is still dominated by solutions for recommending relatively simple and inexpensive products like movies, music, news and books. While there are systems managing more complex item types, such as financial investments or travel, these item categories are considered as atypical cases. Inevitably, complex domains require more elaborated solutions such as those based on knowledge and largely discussed in Chap. 5. Complex products are typically configurable or offered in several variants. This feature still poses a challenge to recommender systems, which are instead designed to consider different configurations as different items. Identifying the more suitable configuration

requires reasoning between the interactions of alternative configurations (classifying and grouping items) and calls for addressing the specificity of the human decision making task generated by the selection of a configuration. In general terms, addressing new types of recommendation domains can call for the introduction of many new and interesting research lines. For instance, Chap. 16 clearly shows how different the recommendation technique must be in domains where reciprocal recommendations are needed, as in dating applications.

As we already indicated in the first edition of this handbook, recommenders that optimize a sequence of recommendations, e.g., a new book every week, are not frequent and we believe that this is still an open issue. It is important to study the sequential dimension of users' decision making both within a recommendation session and between recommendation sessions. Here, we want to further note the importance of such a topic in group recommenders (see Chap. 22). In these systems, sequential recommendations are a natural setting, since stable groups, such as friends or families, repeatedly choose items of the same type, e.g., when deciding where to go for vacation or what to eat at home. A lot more research is needed on algorithms and user interfaces for producing coherent sequences of recommendations. In particular, one should model the effect on users of several contextual conditions such as the manner in which already-shown-items could influence the user evaluation of the next recommendations, or the social role and relationships of the group members.

As it is discussed in Chap. 14 most of the popular recommender systems are now accessed through mobile systems that follow their owners throughout their daily life, and are always within an arm's reach of their owners. In this scenario, as for instance in Google Now, recommender systems can proactively send notifications to their users about items of potential interest that are relevant because of the contextual situation of the user. The challenge is finding true relevant items for the user situation and not overburdening them with a stream of irrelevant interruptions. To address this goal, we must better exploit implicit feedback derived from user usage of the recommendations, but also learn to better identify contextual situations that require push recommendations. We believe that this depends on the detection of contextual changes that are significant to the user and therefore justify a recommendation. For instance, when it is the ideal time for a pause in writing a paper, i.e., the context is changing from work to leisure, a recommendation of a relevant, or personalized, article of sports news can be delivered. Understanding when context changes or could be forced to change and when a user may be receptive to a recommendation push is a challenging issue for further research. As it is also suggested in Chap. 6, in order to develop these new and compelling context aware systems, we need to explore novel engineering solutions to CARS, including: novel data structures, storage systems, user interface components and service oriented architectures.

Another task that we believe should be explored further is guided navigation. This refers to combining classical recommendation lists with tools that let the user navigate more autonomously in the space of possible options. User action interpretation refers to the possibility that in addition to explicit ratings, there could

be many more actions performed by the user operating the recommender that can be detected, analyzed and used to build a better prediction model. The idea is that every single user action should be exploited in the recommendation process. But it is challenging to interpret the user's actions, i.e., the intent behind an action, and there are actions that should be discarded because they were not produced by genuine users, such as actions performed by different users on the same browser, or false and malicious registrations or data or log data caused by robots or crawlers.

Finally, we want to note again that in order to deliver a small set of relevant recommendations to the user, correctly predicting the user rating is only one option. An alternative consists of predicting how the user would compare or rank the available options. This is nowadays an important line of research in recommender systems. Chapter 11 discusses this issue in the context of the Netflix recommender.

Finally, we hope that this handbook, as a useful tool for practitioners and researchers, will contribute to further developing knowledge in this exciting and useful research area and provide a baseline for further exploring the above mentioned issues. Currently the research on RSs has greatly benefited from the combined interest and efforts that industry and academia have invested in this field. We therefore wish the best to both groups as they read this handbook and we hope that it will attract even more researchers to work in this highly interesting and challenging field.

References

1. Adomavicius, G., Tuzhilin, A.: Personalization technologies: a process-oriented perspective. *Commun. ACM* **48**(10), 83–90 (2005)
2. Adomavicius, G., Tuzhilin, A.: Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* **17**(6), 734–749 (2005)
3. Amatriain, X.: Mining large streams of user data for personalized recommendations. *SIGKDD Explor. Newsl.* **14**(2), 37–48 (2013)
4. Arazy, O., Kumar, N., Shapira, B.: Improving social recommender systems. *IT Professional* **11**(4), 38–44 (2009)
5. Asoh, H., Ono, C., Habu, Y., Takasaki, H., Takenaka, T., Motomura, Y.: An acceptance model of recommender systems based on a large-scale internet survey. In: *Advances in User Modeling - UMAP 2011 Workshops, Girona, Spain, July 11–15, 2011, Revised Selected Papers*, pp. 410–414 (2011)
6. Bailey, R.A.: *Design of comparative experiments*. Cambridge University Press Cambridge (2008)
7. Balabanovic, M., Shoham, Y.: Content-based, collaborative recommendation. *Communication of ACM* **40**(3), 66–72 (1997)
8. Baltrunas, L., Ricci, F.: Experimental evaluation of context-dependent collaborative filtering using item splitting. *User Model. User-Adapt. Interact.* **24**(1–2), 7–34 (2014)
9. Ben-Shimon, D., Tsikinovsky, A., Rokach, L., Meisels, A., Shani, G., Naamani, L.: Recommender system from personal social networks. In: K. Wegrzyn-Wolska, P.S. Szczepaniak (eds.) *AWIC, Advances in Soft Computing*, vol. 43, pp. 47–55. Springer (2007)
10. Berkovsky, S., Kuflik, T., Ricci, F.: Mediation of user models for enhanced personalization in recommender systems. *User Modeling and User-Adapted Interaction* **18**(3), 245–286 (2008)

11. Berkovsky, S., Kuflik, T., Ricci, F.: Cross-representation mediation of user models. *User Modeling and User-Adapted Interaction* **19**(1–2), 35–63 (2009)
12. Billsus, D., Pazzani, M.: Learning probabilistic user models. In: *UM97 Workshop on Machine Learning for User Modeling* (1997). URL <http://www.dfki.de/~bauer/um-ws/>
13. Bobadilla, J., Ortega, F., Hernando, A., Gutierrez, A.: Recommender systems survey. *Knowledge-Based Systems* **46**(0), 109–132 (2013)
14. Borràs, J., Moreno, A., Valls, A.: Intelligent tourism recommender systems: A survey. *Expert Systems with Applications* **41**(16), 7370–7389 (2014)
15. Bridge, D., Göker, M., McGinty, L., Smyth, B.: Case-based recommender systems. *The Knowledge Engineering review* **20**(3), 315–320 (2006)
16. Brusilovsky, P.: Methods and techniques of adaptive hypermedia. *User Modeling and User-Adapted Interaction* **6**(2–3), 87–129 (1996)
17. Burke, R.: Hybrid web recommender systems. In: *The Adaptive Web*, pp. 377–408. Springer Berlin / Heidelberg (2007)
18. Chen, L., de Gemmis, M., Felfernig, A., Lops, P., Ricci, F., Semeraro, G.: Human decision making and recommender systems. *TiiS* **3**(3), 17 (2013)
19. Chen, L., Pu, P.: Critiquing-based recommenders: survey and emerging trends. *User Model. User-Adapt. Interact.* **22**(1–2), 125–150 (2012)
20. Cosley, D., Lam, S.K., Albert, I., Konstant, J.A., Riedl, J.: Is seeing believing? how recommender system interfaces affect users’ opinions. In: *In Proceedings of the CHI 2003 Conference on Human factors in Computing Systems*, pp. 585–592. Fort Lauderdale, FL (2003)
21. Ekstrand, M.D., Harper, F.M., Willemsen, M.C., Konstan, J.A.: User perception of differences in recommender algorithms. In: *Eighth ACM Conference on Recommender Systems, RecSys ’14*, Foster City, Silicon Valley, CA, USA - October 06 - 10, 2014, pp. 161–168 (2014)
22. Fisher, G.: User modeling in human-computer interaction. *User Modeling and User-Adapted Interaction* **11**, 65–86 (2001)
23. Golbeck, J.: Generating predictive movie recommendations from trust in social networks. In: *Trust Management, 4th International Conference, iTrust 2006*, Pisa, Italy, May 16–19, 2006, Proceedings, pp. 93–104 (2006)
24. Goldberg, D., Nichols, D., Oki, B.M., Terry, D.: Using collaborative filtering to weave an information tapestry. *Commun. ACM* **35**(12), 61–70 (1992)
25. Herlocker, J., Konstan, J., Riedl, J.: Explaining collaborative filtering recommendations. In: *In proceedings of ACM 2000 Conference on Computer Supported Cooperative Work*, pp. 241–250 (2000)
26. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. *ACM Transaction on Information Systems* **22**(1), 5–53 (2004)
27. Jannach, D., Zanker, M., Felfernig, A., Friedrich, G.: *Recommender Systems: An Introduction*. Cambridge University Press (2010)
28. Kaminskis, M., Ricci, F.: Contextual music information retrieval and recommendation: State of the art and challenges. *Computer Science Review* **6**(2–3), 89–119 (2012)
29. Knijnenburg, B.P., Willemsen, M.C., Gantner, Z., Soncu, H., Newell, C.: Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction* **22**(4–5), 441–504 (2012). DOI [10.1007/s11257-011-9118-4](https://doi.org/10.1007/s11257-011-9118-4)
30. Konstan, J.A., Riedl, J.: Recommender systems: from algorithms to user experience. *User Modeling and User-Adapted Interaction* **22**(1–2), 101–123 (2012)
31. Koren, Y., Bell, R.M., Volinsky, C.: Matrix factorization techniques for recommender systems. *IEEE Computer* **42**(8), 30–37 (2009)
32. Linden, G., Smith, B., York, J.: Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing* **7**(1), 76–80 (2003)
33. Lops, P., de Gemmis, M., Semeraro, G.: Content-based recommender systems: State of the art and trends. In: F. Ricci, L. Rokach, B. Shapira, P.B. Kantor (eds.) *Recommender Systems Handbook*, pp. 73–105. Springer Verlag (2011)

34. Lu, L., Medo, M., Yeung, C.H., Zhang, Y.C., Zhang, Z.K., Zhou, T.: Recommender systems. *Physics Reports* **519**(1), 1–49 (2012)
35. Mahmood, T., Ricci, F.: Improving recommender systems with adaptive conversational strategies. In: C. Cattuto, G. Ruffo, F. Menczer (eds.) *Hypertext*, pp. 73–82. ACM (2009)
36. Mahmood, T., Ricci, F., Venturini, A.: Improving recommendation effectiveness by adapting the dialogue strategy in online travel planning. *International Journal of Information Technology and Tourism* **11**(4), 285–302 (2009)
37. McFee, B., Bertin-Mahieux, T., Ellis, D.P., Lanckriet, G.R.: The million song dataset challenge. In: *Proceedings of the 21st International Conference Companion on World Wide Web, WWW '12 Companion*, pp. 909–916. ACM, New York, NY, USA (2012)
38. McNee, S.M., Riedl, J., Konstan, J.A.: Being accurate is not enough: how accuracy metrics have hurt recommender systems. In: *CHI '06: CHI '06 extended abstracts on Human factors in computing systems*, pp. 1097–1101. ACM Press, New York, NY, USA (2006)
39. Montaner, M., López, B., de la Rosa, J.L.: A taxonomy of recommender agents on the internet. *Artificial Intelligence Review* **19**(4), 285–330 (2003)
40. Park, D.H., Kim, H.K., Choi, I.Y., Kim, J.K.: A literature review and classification of recommender systems research. *Expert Systems with Applications* **39**(11), 10,059–10,072 (2012)
41. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J.: Grouplens: An open architecture for collaborative filtering of netnews. In: *Proceedings ACM Conference on Computer-Supported Cooperative Work*, pp. 175–186 (1994)
42. Resnick, P., Varian, H.R.: Recommender systems. *Communications of the ACM* **40**(3), 56–58 (1997)
43. Ricci, F.: Travel recommender systems. *IEEE Intelligent Systems* **17**(6), 55–57 (2002)
44. Ricci, F.: Recommender systems: Models and techniques. In: *Encyclopedia of Social Network Analysis and Mining*, pp. 1511–1522. Springer (2014)
45. Ricci, F., Cavada, D., Mirzadeh, N., Venturini, A.: Case-based travel recommendations. In: D.R. Fesenmaier, K. Woeber, H. Werthner (eds.) *Destination Recommendation Systems: Behavioural Foundations and Applications*, pp. 67–93. CABI (2006)
46. Robillard, M.P., Maalej, W., Walker, R.J., Zimmermann, T. (eds.): *Recommendation Systems in Software Engineering*. Springer (2014)
47. Schafer, J.B., Frankowski, D., Herlocker, J., Sen, S.: Collaborative filtering recommender systems. In: *The Adaptive Web*, pp. 291–324. Springer Berlin / Heidelberg (2007)
48. Schafer, J.B., Konstan, J.A., Riedl, J.: E-commerce recommendation applications. *Data Mining and Knowledge Discovery* **5**(1/2), 115–153 (2001)
49. Schwartz, B.: *The Paradox of Choice*. ECCO, New York (2004)
50. van Setten, M., McNee, S.M., Konstan, J.A.: Beyond personalization: the next stage of recommender systems research. In: R.S. Amant, J. Riedl, A. Jameson (eds.) *IUI*, p. 8. ACM (2005)
51. Shardanand, U., Maes, P.: Social information filtering: algorithms for automating “word of mouth”. In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI'95)*, pp. 210–217 (1995)
52. Sinha, R.R., Swearingen, K.: Comparing recommendations made by online systems and friends. In: *DELOS Workshop: Personalisation and Recommender Systems in Digital Libraries* (2001)
53. Swearingen, K., Sinha, R.: Beyond algorithms: An HCI perspective on recommender systems. In: J.L. Herlocker (ed.) *Recommender Systems, papers from the 2001 ACM SIGIR Workshop*. New Orleans, LA - USA (2001)
54. Taghipour, N., Kardan, A., Ghidary, S.S.: Usage-based web recommendations: a reinforcement learning approach. In: *Proceedings of the 2007 ACM Conference on Recommender Systems, RecSys 2007, Minneapolis, MN, USA, October 19–20, 2007*, pp. 113–120 (2007)
55. Verbert, K., Drachler, H., Manouselis, N., Wolpers, M., Vuorikari, R., Duval, E.: Dataset-driven research for improving recommender systems for learning. In: *Proceedings of the 1st International Conference on Learning Analytics and Knowledge, LAK '11*, pp. 44–53. ACM, New York, NY, USA (2011)