A Taxonomy for Generating Explanations in Recommender Systems

Gerhard Friedrich and Markus Zanker

■ *In recommender systems, explanations serve as an additional type of information that can help users to better understand the system's output and promote objectives such as trust, confidence in decision making, or utility. This article proposes a taxonomy to categorize and review the research in the area of explanations. It provides a unified view on the different recommendation paradigms, allowing similarities and differences to be clearly identified. Finally, the authors present their view on open research issues and opportunities for future work on this topic.*

Typical explanations for product recommendations include phrases such as "the digital camera *Profishot* is well suited to your needs because you would like to take pictures of your children playing soccer" or "it is a lightweight compact camera especially designed for action photos" or, in the movie domain, "the film *Beyond-SF* was extremely well received by science fiction fans and so you will probably enjoy it too." Such information is commonly exchanged between a sales assistant and a customer during in-store recommendation processes and is usually termed an *explanation* (Brewer, Chinn, and Samarapungavan 1998).

We define explanations in recommender systems by two properties. First, they are information about recommendations, where a recommendation is typically a ranked list of items. Second, explanations support objectives defined by the recommender system designer. Tintarev and Masthoff (2011) provide an enumeration of potential objectives that can be intended by a recommender system's explanations.

For example, the intention behind disclosing the reasoning process of the system could be to increase the user's confidence in making the right decision or to provide additional information such that the user can validate the rationality of the proposed purchase. Indeed, explanations can differ significantly; for instance, they may attempt to maximize the user's confidence throughout the shopping experience while ignoring sat-

  ISSN 0738-4602

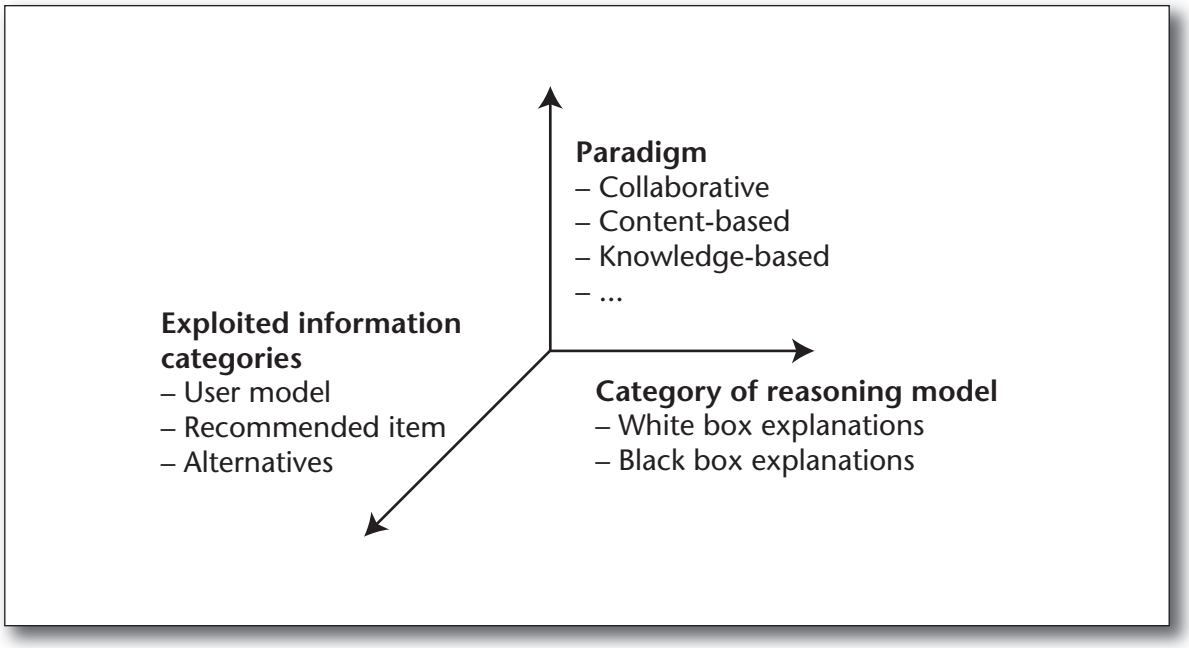

Figure 1. Dimensions for Categorizing Explanation Approaches.

isfaction after purchase. Thus, the possible objectives of explanations are manifold, including aims such as increasing trustworthiness, effectiveness, or persuasiveness, just to name a few.

This article's contribution is a taxonomy that structures the abundance of work that has been produced in the area. We will therefore discuss explanations from different viewpoints and provide an overview of opportunities for future research.

Taxonomy for Generating Explanations

In the following section we will categorize different approaches for explaining recommendations based on major design principles. Figure 1 sketches the three dimensions of the taxonomy, namely the *reasoning model*, the *recommendation paradigm,* and the *exploited information categories*, graphically.

Note that these three dimensions do not exhaustively capture all of the factors that determine the generation of explanations for recommender systems. For instance, the argumentation traces of explanations can be analyzed and classified as has been researched in the domain of expert systems by Ye and Johnson (1995). However, complex argumentation chains that support each claim by justifying and backing statements are typically based on nonmonotonic logical formalisms (Chesñevar, Maguitman, and Loui 2000) whose applicability in highly interactive online environments requires further investigation. In traditional decision support systems (Gönül, Önkal-Atay, and Lawrence 2006) the presentation format and the provision mechanism are also used to classify explanations. Text, images, video, and combinations thereof can be used to explain a system's output. The provision mechanism describes who initiates the generation of explanations. In recommender systems, explanations can be either explicitly requested by users (that is, user invoked) or automatically displayed by the system. Furthermore, structural characteristics such as length, writing style (for instance, the system could use flattering or more factual phrases), or the confidence that is conveyed in explanations can be used as additional dimensions. However, we restrict our taxonomy for generating explanations to those dimensions that are already in use in the literature of the field.

In contrast to classification by design principles, explanations can be categorized by their effects or impact on their users (Tintarev and Masthoff 2011). The impact of explanations on users is measured in terms of the achievement of different objectives that are therefore endogenous constructs (in contrast to the aforementioned exogenous design principles). However, the influence of different explanation strategies with respect to the achievement of objectives such as user satisfaction or trust has not yet been systematically researched. Nevertheless, one can classify approaches for generating explanations by their *intended* goals.

Categories of Reasoning Model for Generating Explanations

The reasoning model describes how explanations are generated and is the most fundamental distinctive criterion. Conventionally, the notion of "explaining recommendations" indicates that the system makes its reasoning process transparent to its user. Looking back in AI's history, expert systems such as the well-known MYCIN (Buchanan and Shortliffe 1984) differentiated between *how* and *why explanations*. In cases where the expert system requires input from the user, explanations as to why this information is needed are provided. Conversely, when an expert system outputs a solution, a client might ask for a justification of the proposed solution. Such explanations are called how explanations because classical expert systems exploited information as to how a conclusion was deduced, for example, the sequence of rules activated in the decision-making process. The implicit assumption was that the model employed to give recommendations can also serve as a model for the argumentation that the advice is plausible from the client's viewpoint.

However, such traces of deduction are not always accepted as high-quality explanations, typically for two reasons. First, depending on the knowledge base, the deduction traces may be far too complex and may confuse the client. This motivated research into high-level explanations (Sqalli and Freuder 1996). Second, the rules exploited for deduction are not necessarily accepted as valid arguments for approving a solution. This argument is based on the observation that the actual production rules in various expert systems only shallowly reflected the known principles and laws of the underlying domain. For example, experience-based rules were formulated for linking symptoms with faults although physical laws exist that could be employed both for drawing conclusions and for justifying the solution based on first principles. The essence of this discussion about deep versus shallow knowledge (Chandrasekara 1991) is that it is the content of knowledge bases and their empirical grounding that are important and not syntactical representation like rules or constraints.

Analogously, similar observations can be made for recommender systems: although the various recommender system approaches agree on the task, the methods exploit different types of information and employ different deduction principles for generating recommendations, that is, collaborative versus content-based versus knowledge-based just to mention the three most important archetypes. For instance, collaborative approaches exploit the similarities between users and items. Knowledge-based recommenders exploit known dependencies between properties of items and users. For instance, someone with a large family and high income will prefer large houses. Therefore, an explanation approach that is based on the exploitation of the recommender's reasoning model must be different for the above-mentioned classes of recommendation systems. However, most of the recently proposed recommendation algorithms that are, for instance, based on matrix factorization models or ensemble learning are unable to explain their recommendations to their users. Herlocker, Konstan, and Riedl (2000) therefore differentiate between white-box and black-box explanations. White-box explanations disclose and exploit the underlying conceptual model of the recommendation engine, while black-box explanations do not disclose the functioning of the system to the user. Reasons for the latter include that the process of computation should not be revealed or there is no sensible way to convey complex computations. For instance, Vig, Sen, and Riedl (2009) use tags to explain that the recommended item and the user's profile are semantically close while Zanker and Ninaus (2010) dynamically combine canned text to generate knowledgable explanations, although the recommendations have been identified and ranked with a different mechanism. To summarize, black-box explanation approaches compute justifications that argue why a specific recommendation is plausible or should be of interest even though the reasoning model did not consider these propositions when actually computing the recommendation. In contrast, white-box explanations disclose (at least partially) the reasoning model and its content to the user in the tradition of expert system's how explanations.

Recommendation Paradigms

Collaborative and content-based filtering as well as knowledge-based recommendations constitute the three basic recommendation paradigms[1] (Jannach et al. 2010). Despite their differences in computing the recommendations, they share the same output function (that is, a ranked list of items) and some ontological commitment that helps to unveil commonalities and generate a unifying view. All personalized recommender systems employ the concepts of *items* and *users*. These concepts, which are effectively sets of entities, are related by predicates (sometimes called associations or roles), such as *likes*(*Users*, *Items*) or *recommendedTo*(*Items*, *Users*). In addition, properties can be exploited to characterize concepts like items and users. These may also be categorized into subconcepts, for example, *price*, *tags*, *degreeOfSatisfaction,* or other sets of feature values. These properties are then employed in relations to describe a domain.

Figure 2a illustrates the domain of collaborative filtering, where preference relations between users

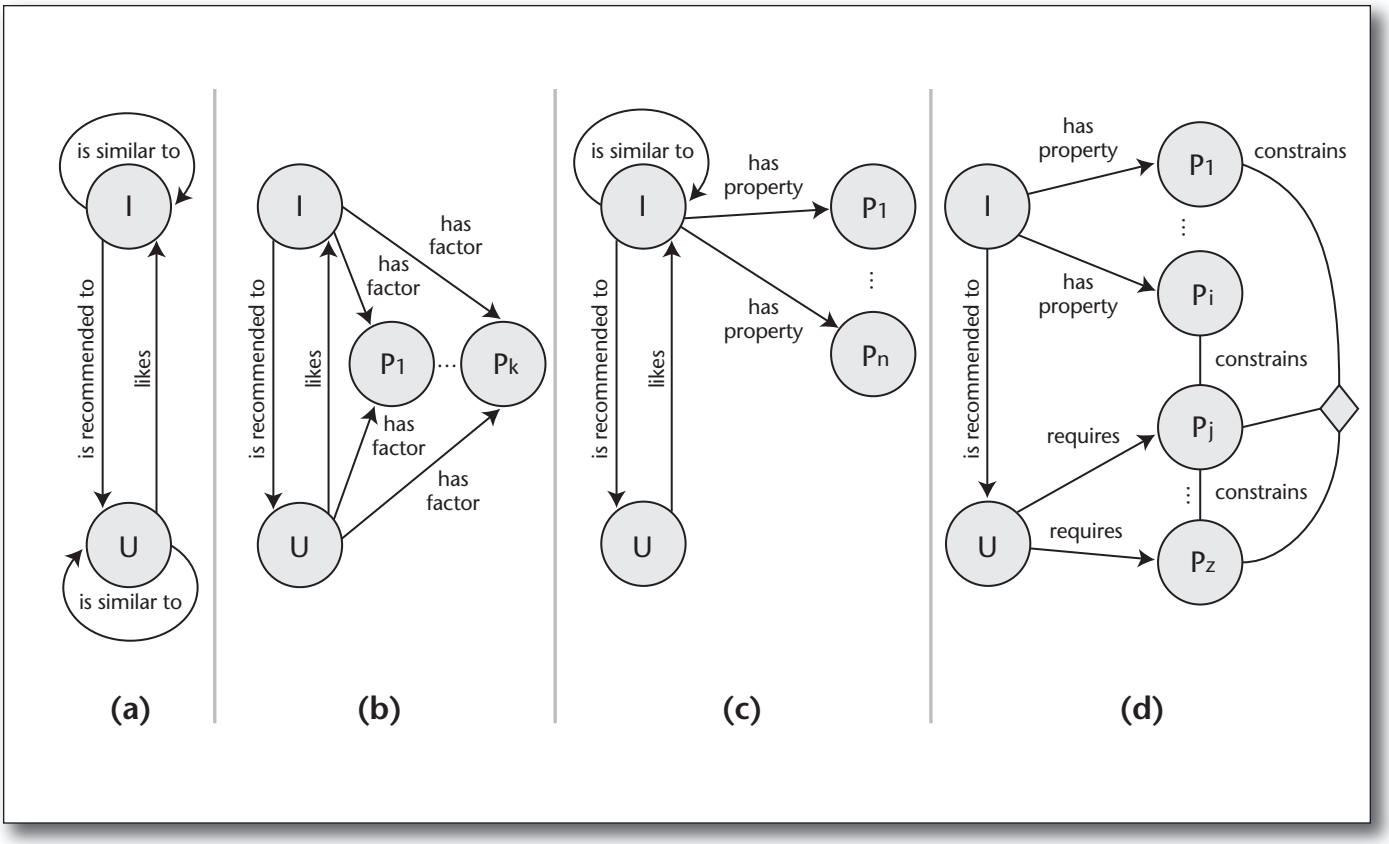

Figure 2. Archetypes of Domain Knowledge.

and items are known (for example, encoded by a rating matrix). Neighborhood-based collaborative filtering algorithms can extend such models by computing either similarity relationships between users or items. In the recommendation phase these similarity relationships are exploited to recommend items that are either similar to the items a user has liked in the past or that are liked by similar users. Specific variants of collaborative filtering, known for their very accurate rating predictions, are matrix factorization (MF) models (Koren, Bell, and Volinsky 2009). The basic model of matrix factorization identifies k abstract factors that can explain most of the signal in the ratings. Figure 2b depicts these factors as abstract property classes $P_1 \ldots P_k$. Explanations of collaborative models disclose similarity relationships between concept instances. However, abstract properties, such as those identified by MF recommenders, rarely capture semantic meaning that can easily be communicated to users. Nevertheless, for the sake of transparency, an MF recommender might compute similarity relationships between users and/or items based on its factor space representation.

The archetype of a content-based recommender system is depicted in figure 2c. There, knowledge about items is given, for instance, in the form of a product catalog with item descriptions or by a term frequency vector. Recommendations are computed by determining those items that are most similar to items the user is already known to like. Content-based white-box explanations therefore explain similarities between items by disclosing their property relationships. The third paradigm encompasses knowledge-based recommendation approaches. They are characterized by additional domain properties such as abstract user requirements or preferences as well as various relationships between them. Knowledge-based recommender systems can encode explicit sales expertise such as which item features help to fulfill a specific user requirement. For instance, a constraint-based recommender system models items and users using sets of variables (represented by $P_1 \ldots P_i$ and $P_j \ldots P_z$ in figure 2d). N-ary relationships between these properties constrain the allowed combinations. Solutions are computed by identifying those items that satisfy all the domain restrictions. In cases where no solution exists, either the user is asked to revise her or his preferences or the reasoner relaxes some of the domain restrictions. See Zanker, Jessenitschnig, and Schmid (2010) for a discussion

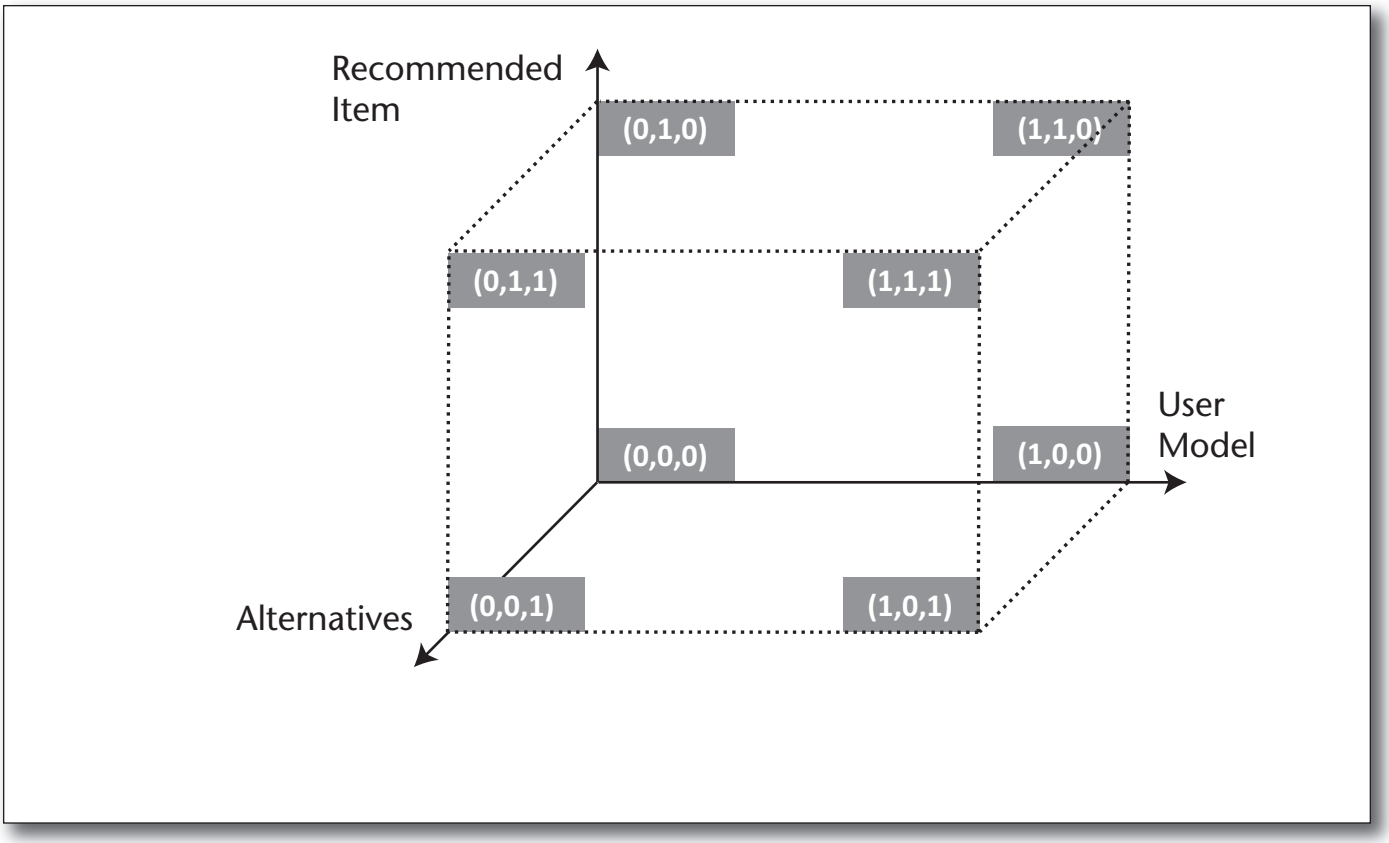

Figure 3. Information Categories.

of different strategies for preference reasoning in constraint-based recommender system. Consequently, explanations disclose variable assignments (such as *has property* and *requires* relationships) and the constraining relationships between them to explain a recommendation (Friedrich 2004). Utility-based recommender systems that, for instance, implement multiattribute utility theory (Felfernig et al. 2008) also fall into this category.

Consequently, the unifying view is that explanation generation (like the computation of recommendations itself in the case of a white-box approach) exploits the relations between users, items, and properties. However, how the instances of these relations are derived depends on the employed paradigm. For instance, the fact that two items are similar can be deduced using most of the aforementioned techniques, but the explanation for such a specific similarity relationship and the associated reasoning principles differ depending on the applied reasoning paradigm.

Information Categories

The third characterizing dimension comprises the information categories that can be exploited for generating explanations. We differentiate between three different aspects of input: user model, recommended item, and alternatives.

User model. Are explanations tailored to the system's beliefs about the given user? For instance, the system could present arguments based on the user's known ratings, preferences, or demographics.

Recommended item. Is the explanation dependent on the specific recommended item? For instance, does the explanation make statements about the specific characteristics of the recommended item?

Alternatives. Do explanations argue in favor of or against alternatives to the recommended item?

Obviously, explanations can be tailored to all three aspects, none of them, or any other combination as depicted in figure 3. For instance, explanations located at position (0,0,0) do not adapt to any of the three information categories and can be realized using static text, for example, explanatory phrases such as "the system made correct predictions in more than 80 percent of all sessions" or "users with page views similar to yours purchased the following products."

Examples that address only the user model in generated explanations (position (1,0,0)) include

"we considered the rating profiles of the 48 most similar peers from Chicago" or "your average rating is 3.5 stars." In contrast, explanations that refer only to known data on the recommended item could be as follows: "the proposed movie won three Golden Globe and two Academy awards" or "this wine is the perfect accompaniment for desserts." However, most commonly one expects explanations to exploit both the user model and characteristics of the recommended item such as a histogram that depicts "your neighbors' ratings for this movie" or "this digital camera is very reliable and is the perfect companion for going abroad; however, it costs a bit more than you specified." The third aspect enables explanations to compare a recommended item with alternative options and argue with their differences. For instance, Pu and Chen (2007) researched designs providing explanation such as "although this is the most popular product, we would also recommend the following products that are cheaper and lighter, but have lower processor speeds." Finally, an explanation that encompasses all three aspects together might include the following: "We think you will enjoy the Carribean Spa Resort, as it offers, for instance, romantic candlelight dinners that are perfect for young couples. Furthermore, this resort is closer to your home and less costly than the second ranked offer."

Commonalities of Explanations

An important principle of explanations is succinctness, which has a long history in knowledge-based systems. Information and knowledge that are not relevant for answering a request should be excluded from an explanation. The historic MYCIN expert system also included this feature, disclosing only rules that actually contributed to an answer. In constraint-based systems only those constraints that are needed to generate the value assignments for entailing a recommendation are considered for explanations. In general, generating explanations can be seen as a form of abductive reasoning (Eiter, Gottlob, and Leone 1997), where the goal is to find the smallest subsets of a set of logical sentences such that a query is entailed. For example a recommendation for a car to someone who has many children and does not care about driving performance will be based on the number of seats and the volume of the car. Therefore, arguing based on horsepower is irrelevant and should thus be excluded from an explanation. More formally, given a knowledge base KB and a query Q (also known as a recommended item in this context) that should be explained, the smallest subset KB' that entails Q is the basis for an explanation; that is, $KB' \subseteq KB$, $KB' \models Q$, and for all $KB'' \subset KB'$, $KB'' \not\models Q$ where $\models$ is the usual entailment relation for logical sentences.

Note that several explanations can exist for a recommendation. In addition to succinctness, systems generating explanations must consider previously presented solutions to avoid spurious explanations (Friedrich 2004). This minimal knowledge base serves as a basis for generating explanations but is not one itself because, in the context of recommender systems, explanations are the inputs for building arguments for humans. Some parts of the knowledge base may be necessary for entailment but are obvious for humans. For instance, classifying a van with seven seats as a family car is obvious for most users whereas a recommender system has to deduce this classification from data sheets in cases where this is not given.

Furthermore, this principle of succinctness and abductive reasoning can be applied to all forms of recommender paradigms. First, collaborative and content-based recommender systems are designed to exploit all relations of the employed ontology for their reasoning process. So the set of relations is the smallest required to derive a recommendation by design. However, the set of relation tuples needed to provide a specific recommendation is, in fact, reduced for the generation of explanations in various systems. For instance, in Herlocker, Konstan, and Riedl (2000) explanations of the following type are considered: "As 23 out of 33 people similar to you like this item, we think you will like it too."[2] Note that the set of similar users is reduced to those relevant for the recommendation. In addition, the identities of individuals unknown to the user are not explicitly shown. Instead, aggregated values are presented because one could assume that a greater level of detail would be meaningless and confuse the user.

Exemplary Practical Systems

Finally, we instantiate the proposed taxonomy, by presenting exemplary work in the field. Table 1 summarizes the categorization of different approaches. Herlocker, Konstan, and Riedl's (2000) paper is the first and most influential one (based on citation count) related to this topic. They compared 21 different explanation styles for collaborative filtering with a participants experimental design and asked users about their intent to see the recommended movies. We denote the best and the second best strategy according to measured user intention in table 1. Their paradigm is collaborative because data about the similarity relationships between users is disclosed. However, explanations may present this data with varying levels of detail by considering or ignoring data from the current user or the recommended item. Bilgic and Mooney (2005) compared content-based and collaborative explanation styles, that is, argumentation based on key words describing the recommended item versus disclosing the behavior of the most similar

Reference	Explanation Style/System	Reasoning Model	Paradigm	Information Categories*
Herlocker et al. (2000)	Histogram with grouping	White box	Collaborative	(1,1,0)
Herlocker et al. (2000)	Past performance	White box	Collaborative	(1,0,0)
Jannach (2004)	Advisor Suite	White box	Knowledge-Based	(1,1,0)
Bilgic and Mooney (2005)	Keyword style	White box	Content-Based	(1,1,0)
Bilgic and Mooney (2005)	Neighbor style	White box	Collaborative	(1,1,0)
Bilgic and Mooney (2005)	Influence style	White box	Collaborative	(1,1,0)
Pu and Chen (2007)	Trade-off explanations	White box	Knowledge-Based	(0,1,1)
Symeonidis et al. (2008)	MoviExplain	White box	Content-Based	(1,1,0)
Felfernig et al. (2009)	Plausible repairs	White box	Knowledge-Based	(1,0,0)
Vig et al. (2009)	Tagsplanations	Black box	Content-Based	(1,1,0)
Zanker and Ninaus (2010)	Knowledgeable explanation	Black box	Knowledge-Based	(1,1,0)

Table 1. Categorization of Explanation Approaches.

*(User model, Recommended item, Alternatives)

peers. Their evaluation measured how informative the explanations appeared to users by comparing user ratings for the recommended items after having read the explanations with reratings for the same items after having experienced them. Not astoundingly, the key-word style performed best in this setting.

Symeonidis, Nanopoulos, and Manolopoulos (2008) further developed the key word and influence explanation styles by establishing a collaborative and content-based hybrid that justifies recommendations as follows: ''Item *A* is suggested because it contains features *i* and *j* that are also included in items *X, Y,* and *Z* that you like." Consequently, by disclosing the similarities between items based on their properties the generated explanations follow a content-based argumentation approach that adapts to what is known about the current user and the recommended item. A somewhat similar but more encompassing idea has been proposed by Vig, Sen, and Riedl (2009), who select the presumably most influential (user generated) tags instead of key words to justify recommendations. Furthermore, they apply this explanation technique as a black-box approach where state-of-the-art collaborative filtering mechanisms compute the recommendations themselves.

In the category of knowledge-based recommendation systems, the Advisor Suite system (Jannach 2004) is an example of a system that discloses the sets of successfully applied and relaxed filter constraints to its users. Felfernig et al. (2009) extend such conversational recommendation systems by proposing a mechanism to compute repair actions that explain to users how to modify their specified requirements in order to avoid empty recommendation lists. In Zanker and Ninaus (2010) the generation of explanations from declarative knowledge bases is generalized to a predicate-based finite state automata that is meant for the sole purpose of flexibly composing explanations from canned text.

Finally, trade-off explanations provide additional information about alternatives by comparing the features of a recommended item to the characteristics of other items, that is, addressing the information category *alternatives* in our taxonomy. Pu and Chen (2007) developed a trust model and reported on the potential of these explanation interfaces in building trust relationships between such a system and its users.

Open Research Issues

Currently most implementations of explanation components in recommender systems follow a white-box approach. This is reasonable in order to optimize development effort. However, from a scientific point of view little research concentrates on predicting which explanation strategies are best suited for achieving which explanation objectives. Consequently, we anticipate substantial improvement in the quality of explanations, provided that researchers develop a better understanding of how different explanation traits affect users in order to develop an effective explanation model for particular domains.

Basically, we see two lines of research toward more effective explanations: first, the creation of new kinds of information, interaction, and pres-

entation styles that can be exploited for generating explanations, and second, analysis of how this information, interaction, and presentation affect various explanation objectives and under which conditions. For example, the development of a theory that classifies which types of users, situations, items, and explanation objectives favor a particular explanation model.

The development of various kinds of explanations is affected by the current technical state of the art. For instance, social networks and the increased provision of information by customers (for example, tags, reviews and other forms of user-generated content, geospatial data, friendship relations, linked data, and others) will drive future research toward novel data mash-ups and interaction scenarios. Thus, web 2.0 and the semantic web will enable many new forms of explanations.

In terms of theory, we expect that insights from related disciplines like psychology and marketing will help to explore systematically the impact and trade-offs of different aspects of explanation strategies. Psychology has elaborated a rich body of knowledge that can be exploited in recommender systems to achieve explanation objectives. For instance, framing and formulating decision problems in different ways produces predictable shifts of preference (Tversky and Kahneman 1981). As explanations (at least partially) help to construct a user's decision frame, they need to be carefully designed to consider these effects. In Gönül, Önkal-Atay, and Lawrence (2006), a significant interaction effect between the length and the conveyed confidence in explanations of decision support systems is reported. One of the article's findings is that users appreciate long textual explanations only if the system conveys confidence. See Teppan and Felfernig (2009) and Yoo and Gretzel (2011) for further discussions on the application of different psychological effects in recommender systems leading to persuasive system traits.

Furthermore, if there is a clear understanding of the conditions under which a specific explanation form is most effective, it is reasonable to design methods that increase the likelihood of generating the right explanation by reducing the uncertainty about users' mental states. For example, if we know that a certain type of customer favors specific types of arguments, methods to classify the customer are of increased importance.

Although it is desirable to develop a perfect theory that can predict the most appropriate explanations for all contexts and situations of online recommendations, we anticipate that developers of recommendation systems still have to accept some uncertainty with respect to the quality and impact of their designed explanation strategies. Consequently, the question is how observations from interactions with conversational recommender systems can be exploited to improve explanations and subsequently the overall performance of the recommender system itself. One initial piece of work in this direction is reported in Mahmood, Ricci, and Venturini (2009). There different conversational moves of a recommender system (and different explanation strategies definitely constitute such moves) are interpreted as actions in a probabilistic state model, and the system's performance can be optimized with the aim of reaching some preferred goal states.

To sum up, there is plenty of evidence that explanations are a powerful tool for influencing user behavior. With the increased importance of recommender systems in the sales process, interest in explanations is increasing. AI methods and web 2.0 technologies combined with insights from neighboring fields have already developed several successful explanation approaches. However, given that we have only exploited a small fraction of the available knowledge about human behavior, we anticipate many interesting and valuable research results related to this topic in the near future.

Notes

1. Note that other authors use slightly different groupings, such as collaborative, content-based, and hybrid (Adomavicius and Tuzhilin 2005).

2. This is a textual description of a graphical representation contained in Herlocker, Konstan, and Riedl (2000).

References

Adomavicius, G., and Tuzhilin, A. 2005. 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.

Bilgic, M., and Mooney, R. J. 2005. Explaining Recommendations: Satisfaction Versus Promotion. Paper presented at Beyond Personalization 2005: A Workshop on the Next Stage of Recommender Systems Research, San Diego, CA, 9–12 Jan.

Brewer, W. F.; Chinn, C. A.; and Samarapungavan, A. 1998. Explanation in Scientists and Children. *Minds and Machines* 8(1): 119–136.

Buchanan, B., and Shortliffe, E. 1984. *Rule Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley.

Chandrasekara, B. 1991. Models Versus Rules, Deep Versus Compiled Content Versus Form: Some Distinctions in Knowledge Systems Research. *IEEE Expert* 6(2): 75–79.

Chesñevar, C. I.; Maguitman, A. G.; and Loui, R. P. 2000. Logical Models of Argument. *ACM Computing Surveys* 32(4): 337–383.

Eiter, T.; Gottlob, G.; and Leone, N. 1997. Abduction from Logic Programs: Semantics and Complexity. *Theoretical Computer Science* 189(1–2): 129–177.

Felfernig, A.; Friedrich, G.; Schubert, M.; Mandl, M.; Mairitsch, M.; and Teppan, E. 2009. Plausible Repairs for Inconsistent Requirements. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, 791–796. Menlo Park, CA: AAAI Press.

Felfernig, A.; Friedrich, G.; Teppan, E.; and Isak, K. 2008. Intelligent Debugging and Repair of Utility Constraint Sets in Knowledge-Based Recommender Applications. In *Proceedings of the 13th ACM International Conference on Intelligent User Interfaces*, 218–226. New York: Association for Computing Machinery.

Friedrich, G. 2004. Elimination of Spurious Explanations. In *Proceedings of the 16th European Conference on Artificial Intelligence*, ed. R. L. de Mantaras and L. Saitta, 813–817. Amsterdam, The Netherlands: IOS Press.

Gönül, M. S.; Önkal-Atay, D.; and Lawrence, M. 2006. The Effects of Structural Characteristics of Explanations on Use of a DSS. *Decision Support Systems* 42(3): 1481–1493.

Herlocker, J. L.; Konstan, J. A.; and Riedl, J. 2000. Explaining Collaborative Filtering Recommendations. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, 241–250. New York: Association for Computing Machinery.

Jannach, D. 2004. Advisor Suite—A Knowledge-Based Sales Advisory System. In *Proceedings of the 16th European Conference on Artificial Intelligence*, ed. R. L. de Mantaras and L. Saitta, 720–724. Amsterdam, The Netherlands: IOS Press.

Jannach, D.; Zanker, M.; Felfernig, A.; and Friedrich, G. 2010. *Recommender Systems: An Introduction*. New York: Cambridge University Press.

Koren, Y.; Bell, R. M.; and Volinsky, C. 2009. Matrix Factorization Techniques for Recommender Systems. *IEEE Computer* 42(8): 30–37.

Mahmood, T.; Ricci, F.; and Venturini, A. 2009. Improving Recommendation Effectiveness: Adapting a Dialogue Strategy in Online Travel Planning. *International Journal of Information Technology and Tourism* 11(4): 285–302.

Pu, P., and Chen, L. 2007. Trust-Inspiring Explanation Interfaces for Recommender Systems. *Knowledge-Based Systems* 20(6): 542–556.

Ricci, F.; Rokach, L.; Shapira, B.; and Kantor, P. B., eds. 2011. *Recommender Systems Handbook*. Berlin: Springer.

Sqalli, M. H., and Freuder, E. C. 1996. Inference-Based Constraint Satisfaction Supports Explanation. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference*, 318–325. Menlo Park, CA: AAAI Press.

Symeonidis, P.; Nanopoulos, A.; and Manolopoulos, Y. 2008. Providing Justifications in Recommender Systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part A* 38(6): 1–1272.

Teppan, E. C., and Felfernig, A. 2009. Asymmetric Dominance- and Compromise Effects in the Financial Services Domain. In *Proceedings of the 2009 IEEE Conference on Commerce and Enterprise Computing,* ed. B. Hofreiter and H. Werthner, 57–64. Los Alamitos, CA: IEEE Computer Society.

Tintarev, N., and Masthoff, J. 2011. Designing and Evaluating Explanations for Recommender Systems. In *Recommender Systems Handbook,* ed. F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, 479–510. Berlin: Springer.

Tversky, A., and Kahneman, D. 1981. The Framing of Decisions and the Psychology of Choice. *Science* 211(4481): 453–458.

Vig, J.; Sen, S.; and Riedl, J. 2009. Tagsplanations: Explaining Recommendations Using Tags. In *Proceedings of the 14th International Conference on Intelligent User Interfaces,* 47–56. New York: Association for Computing Machinery.

Ye, L. R., and Johnson, P. E. 1995. The Impact of Explanation Facilities in User Acceptance of Expert System Advice. *MIS Quarterly* 19(2): 157–172.

Yoo, K. H., and Gretzel, U. 2011. Creating More Credible and Persuasive Recommender Systems: The Influence of Source Characteristics on Recommender System Evaluations. In *Recommender Systems Handbook,* ed. F. Ricci, L. Rokach, B. Shapira, and P. B. Kantor, 455–477. Berlin: Springer.

Zanker, M., and Ninaus, D. 2010. Knowledgeable Explanations for Recommender Systems. In *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, 657–660. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Zanker, M.; Jessenitschnig, M.; and Schmid, W. 2010. Preference Reasoning with Soft Constraints in Constraint-Based Recommender Systems. *Constraints* 15(4): 574–595.

Gerhard Friedrich is a full professor at the Alpen-Adria University (AAU), Klagenfurt, Austria. He received his PhD at the Technical University Vienna, Austria. Currently he is the director of the institute of applied informatics at AAU and leads the research group on intelligent systems and business informatics. He coauthored *Recommender Systems: An Introduction* and has published numerous papers in the area of artificial intelligence focusing on recommender systems, diagnosis, knowledge acquisition, and configuration. Friedrich is an area chair of *AI Communications*.

Markus Zanker is an associate professor at the Alpen-Adria University, Klagenfurt, Austria. He received his PhD at the AAU and directs the research group on recommender systems. In 2010 he was the program cochair of the 4th International ACM Conference on Recommender Systems. He coauthored *Recommender Systems: An Introduction* and has published numerous papers in the area of artificial intelligence focusing on recommender systems, consumer buying behavior, and human factors. Zanker is an associate editor of the *International Journal of Human-Computer Studies*.