

## Recommendation to Groups

Anthony Jameson<sup>1</sup> and Barry Smyth<sup>2</sup>

<sup>1</sup> DFKI, German Research Center for Artificial Intelligence

<sup>2</sup> Department of Computer Science, University College Dublin

**Abstract.** Recommender systems have traditionally recommended items to individual users, but there has recently been a proliferation of recommenders that address their recommendations to groups of users. The shift of focus from an individual to a group makes more of a difference than one might at first expect. This chapter discusses the most important new issues that arise, organizing them in terms of four subtasks that can or must be dealt with by a group recommender: 1. acquiring information about the user's preferences; 2. generating recommendations; 3. explaining recommendations; and 4. helping users to settle on a final decision. For each issue, we discuss how it has been dealt with in existing group recommender systems and what open questions call for further research.

### 20.1 Introduction

Almost all of the techniques of web-based personalization discussed in the other chapters of this book are designed to allow effective adaptation to individual users. But often the users of such systems operate not individually but in groups, which may vary from formally established, long-term groups to ad hoc collections of individuals who use a system together on a particular occasion. This phenomenon can in principle occur with just about any form of web personalization. In this chapter, we will focus on the subclass of recommender systems (cf. the chapters in this volume by Schafer et al. [34]; Pazzani & Billsus [31]; Smyth [35]; Burke [5]; and Goy & Ardissono [15]), but many of the points made will be applicable by analogy to other types of adaptive web-based system (cf. Section 20.6.2).

Some types of items that a system can recommend (e.g., restaurants and museum exhibits; see Table 20.1 for additional examples) tend to be used at least as often by groups as by individuals, so addressing recommendations to individuals can actually be unnatural. Moreover, the evolution of computers away from the desktop PC makes it increasingly natural for systems to address groups as well as individuals: Wall displays, information kiosks, PDAs, and cell phones can be used easily by persons who are interacting with each other. And even with the traditional PC, users are being offered an increasing variety of ways to communicate with each other and perform tasks together. For these reasons, we can expect a continuing growth in the trend toward recommendation (and, more generally, adaptation) to groups of users.

In this chapter, we will identify the issues that should be addressed by designers of group recommender systems and the ways in which they have been dealt with in

systems that have been developed so far. Our two goals are to allow designers and researchers (a) to make effective use of knowledge and experience that have already been accumulated and (b) to address the many open questions that still require careful consideration and research.

### **20.1.1 Existing Group Recommenders**

Figure 20.1 lists almost all of the group recommender systems that, according to the authors' knowledge, have been described in the literature up to the time of the writing of this chapter. Although a number of these systems have been described only briefly and some were not implemented as web-based systems, we will refer to aspects of all of them, so as to convey an idea of the variety of application settings, design issues, and possible methods that designers of group recommender systems should be aware of.

### **20.1.2 Overview of Recommendation Subtasks and Issues**

Relative to recommendation for individuals, there are a number of new issues that arise with group recommenders. Table 20.2 organizes these issues in terms of four high-level subtasks that must (or can) be performed by a group recommender. The issues corresponding to each of these subtasks will be addressed in one section of this chapter.

These subtasks differ greatly in the amount of attention they have attracted in research so far and hence in the length of the corresponding sections of this chapter. By far the most research has been done on Subtask 2, a fact that is understandable given that any group recommender must have some way of assessing the suitability of items for the group. Subtask 3 is the second most popular one, especially given the growing interest in making the reasoning underlying recommendations comprehensible to users (cf., e.g., the chapter in this volume by Schäfer et al. [34]). Subtask 1 has attracted much less attention, because many methods for acquiring information about users' preferences are equally applicable to groups and to individuals; though the extension to groups does not always require a change in methods, it can create opportunities to introduce new methods. Subtask 4 has attracted the least attention of all, since it is usually assumed that the final decision about whether to accept a recommendation will be made by a single group member or in face-to-face discussion among group members.

## **20.2 Acquiring Information About Group Members' Preferences**

Most group recommenders developed so far apply methods for acquiring information about users' preferences that are barely distinguishable from the methods applied in recommender systems for individuals. After briefly surveying some typical applications of such methods, we will look at preference acquisition methods that have been developed specifically for group recommendation settings.

**Table 20.1.** Overview of the group recommender systems mentioned in this chapter.

System	Reference	(Examples of) Groups of Users	Items recommended
<i>Web / news pages</i>			
Let's Browse	Lieberman et al. (1999)	Persons browsing the web together	Web pages
G.A.I.N	Pizzutilo et al. (2005)	Persons viewing a wall display or information kiosk	News items
I-Spy	Smyth et al. (2005)	Employees of a company	Web pages
<i>Tourist attractions</i>			
Intrigue	Ardissono et al. (2003)	Tourists	Sightseeing tours
CATS	McCarthy et al. (2006)	Friends planning a vacation	Vacation packages
Travel Decision Forum	Jameson (2004)	Friends planning a vacation	Criteria for choosing a vacation package
Group Modeler	Kay and Niu (2005)	Persons visiting a museum together	Information about exhibits
Pocket RestaurantFinder	McCarthy (2002)	Colleagues going out to dine together	Restaurants
<i>Music tracks</i>			
MusicFX	McCarthy and Anagnost (1998)	Persons working out in a gym	Music stations
Flytrap	Crossen et al. (2002)	Persons using a public area of a building	Music tracks to be played
In-vehicle multimedia recommender	Yu et al. (2005)	Passengers in a vehicle	Multimedia items to be played
Adaptive Radio	Chao et al. (2005)	Colleagues working together in an office	Songs to be played on the radio
<i>Television programs and movies</i>			
FIT	Goren-Bar and Glinansky (2002)	Family members watching TV together	TV programs
TV program recommender	Yu et al. (2006)	TV viewers	Sequences of TV programs
PolyLens	O'Connor et al. (2001)	Persons planning to go to a movie together	Movies

**Table 20.2.** Overview of the issues to be addressed in this chapter, organized in terms of four subtasks of a group recommender system.

Subtask of the recommender system	Difference from recommendation to individuals	General issues raised
1. The system acquires information about the members' preferences.	If members specify their preferences explicitly, it may be desirable for them to be able to examine each other's preference specifications.	What benefits and drawbacks can such examination have, and how can it be supported by the system?
2. The system generates recommendations.	Some procedure for predicting the suitability of items for a group as a whole must be applied.	What conditions might such a procedure be required to fulfill; and what kinds of procedure tend to fulfill these conditions?
3. The system presents recommendations to the members.	The (possibly different) suitability of a solution for the individual members becomes an important aspect of a solution.	How can relevant information about suitability for individual members be presented effectively?
4. The system helps the members arrive at a consensus about which recommendation (if any) to accept.	The final decision is not necessarily made by a single person; negotiation may be required.	How can the system facilitate the necessary communication among group members?

### 20.2.1 Preference Acquisition Methods That Are Not Specifically Adapted to Group Recommendation

**Acquisition of Preferences Without Explicit Specification.** As we have seen in the chapters by Schafer et al. [34]; Pazzani & Billsus [10]; Smyth [35]; Burke [5]; and Goy & Ardissono [15], many recommender systems do not require their users to specify their preferences explicitly. With group recommenders as well, it may be possible for the system to get by with implicitly acquired information about users. A straightforward example is found in the system FLYTRAP (Crossen et al. [11]), which selects music for playing in a public room. The system learns about the music preferences of the potential users by (a) noticing what MP3 files each user plays on his or her own computer and (b) consulting available information about the music played to derive a model of each user's preferences.

Another example is found in LET's BROWSE (Lieberman et al. [23]), which recommends web pages to a group of two or more persons who are browsing the web together. The system makes initial estimates of the interests of its users by analyzing the words that occur in each user's web homepage. During the actual group browsing, it analyzes the words that occur in the pages visited by the group.

**Explicit Preference Specification.** But there are some types of group recommender that do require an explicit specification of preferences. An example is the **POCKET RESTAURANTFINDER** (McCarthy [26]), which helps a group of people who are preparing to go out to eat together in selecting a restaurant. For each of 15 types of cuisine that might be represented at a given restaurant, each user must indicate his or her preference on a 5-point scale ranging from “Definitely don’t want ...” to “Definitely want ...”. Similar ratings are given for 17 possible restaurant amenities, 3 price categories, and 3 ranges of travel time from the current location. (Users presumably consider this rating effort worthwhile only if they intend to use the system repeatedly.)

Similarly, explicit preference specifications are required by the **TRAVEL DECISION FORUM** (Jameson [18]; Jameson et al. [19]), which helps a group of users to agree on the desired attributes of a vacation that they are planning to take together. The system needs to know how each user feels about dozens of attributes of vacation destinations, ranging from the facilities that are available in their rooms to the sightseeing attractions that are available in the surrounding area. Here again, only explicit elicitation is likely to be feasible.

A less explicit form of preference specification is found in **POLYLENS**, an extension of the **MOVIELENS** system (cf. the chapter in this volume by Schäfer et al. [34]) that recommends movies to groups of users. Since **POLYLENS** (like **MOVIELENS**) is based on collaborative filtering, users do not explicitly describe their movie preferences, but they do rate individual movies on a scale from 1 to 5 stars. As we will see, this procedure raises some of the same issues as the explicit specification of general preferences.

An intermediate case between implicit and explicit preference specification is found in the system **I-SPY** (see, e.g., Smyth et al. [37]), a community-based web search engine that personalizes search results for a community of like-minded searchers on the basis of a model of community search preferences.<sup>1</sup> The **I-SPY** user indicates interest in a given search result by selecting the result in question from a query result list, and **I-SPY** interprets each result selection as an indication of relevance with respect to the current search query. The specification is implicit in that the user’s primary intent in selecting a result is not in general to indicate his or her preferences to the system; but it has some elements of explicitness in that users are aware of the fact that their selections are being interpreted as reflecting their preferences and can, if they like, choose results that they would not otherwise have chosen, in order to influence the system’s preference model (cf. the discussion of manipulation in Section 20.3.2).

### 20.2.2 Adapting Preference Specification to the Requirements of Group Recommendation

**Focus on Negative Preferences.** In the context of the system **ADAPTIVE RADIO**, Chao et al. [7] argue that the method used to elicit preferences from users should take into account the way in which these preference specifications will subsequently be used

<sup>1</sup> **I-SPY** is not a group recommender system in the most commonly assumed sense, because its recommendations are made use of by individuals, not by members of a collaborating group. But as we will see, the system addresses some of the same issues as systems that make recommendations to groups.

for the generation of recommendations. Specifically, they argue that for a system that chooses music to be played to a group, it makes more sense to elicit *negative preferences* (e.g., expressions of dissatisfaction with particular music tracks) than to elicit more detailed types of rating such as the ones mentioned in the previous subsection. If the procedure that will be used for generating recommendations is designed mainly to avoid the playing of music that is disliked by any member (cf. Section 20.3.2), effort expended by one group member in discriminating among different degrees of liking may in fact be wasted, since many of the songs that that member likes may in effect be vetoed by another member anyway. An informal evaluation of ADAPTIVE RADIO suggested that the focus on negative preferences was appropriate in the particular application setting studied.

**Sharing Information About Specified Preferences.** In a recommender system for individuals, there is in general no person besides the user who has an immediate interest in seeing explicitly specified preferences with a view to improving the current recommendation process. In a group recommender, each member may have some interest in knowing the other members' preferences, for several possible reasons:

1. *Saving of effort.* Specifying preferences is usually seen by users as a tedious process. (The avoidance of tedium is claimed by Chao et al. [7], as a supplementary advantage of their focus on negative preferences.) If a group member  $m_1$  knows that another member  $m_2$  with generally similar preferences has already specified their preferences,  $m_1$  may be able to save time and effort by copying at least some of  $m_2$ 's entries and then perhaps making some changes—especially if the system makes it easy to do such copying and postediting.
2. *Learning from other members.* Another member's preferences may be based in part on knowledge or experience (e.g., concerning a particular vacation destination) that the current member lacks.

An attempt to exploit both of these potential benefits is found in the TRAVEL DECISION FORUM: A simple extension of a typical rating-scale dialog box allows the current member optionally to view (and perhaps copy) the preferences already specified by other members (see Figure 20.1). An additional feature that makes sense mainly if other persons will be viewing the specifications is the option to add brief verbal explanations or *arguments* for specific ratings.<sup>2</sup> Arguments can have various forms and functions in group decision contexts (cf., e.g., Jennings et al. [21]). In a group recommendation context, two typical functions are (a) to persuade other members to specify a similar preference, perhaps by giving them information that they previously lacked; and (b) to explain and justify a member's preference even if the argument is not generalizable to other members (e.g., "I can't go hiking, because of an injury").

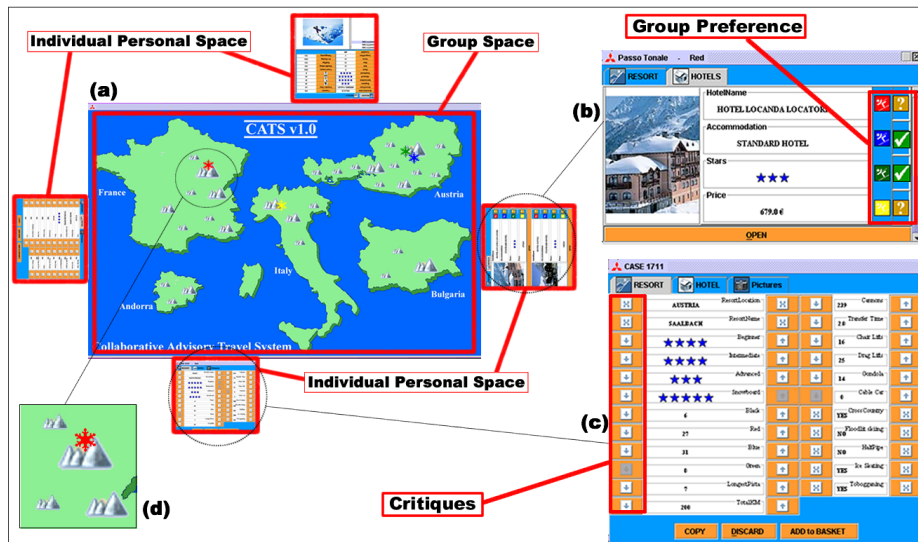
Experience with this method of *collaborative preference specification* has revealed further benefits beyond the two already mentioned:

<sup>2</sup> These arguments can be entered and viewed in pop-up windows that are not visible in Figure 20.1.

**Fig. 20.1.** Dialog box for the collaborative specification of preferences in the TRAVEL DECISION FORUM. (The currently active group member is Claudia, the other two are Ritchie and Tina. The preferences of each member are represented by the first letter of his or her name. Each scale refers to a single attribute and ranges from -- for “Don’t want it” to ++ for “Want it”. The highlighting of one cell for each attribute is added only when a compromise proposal has been suggested, as explained in Section 20.3.1. Figure 1 of Jameson [18], reproduced with permission.)

1. *Taking into account attitudes and anticipated behavior of other members.* Sometimes the preference of the current member depends in part on the preferences and/or the anticipated behavior of one or more other members. For example, if  $m_1$  sees that  $m_2$  has specified a strong preference for tennis facilities,  $m_1$  may want to specify a similar preference, reasoning that if a hotel is found that offers tennis,  $m_1$  and  $m_2$  will be able to play together. Otherwise,  $m_1$  may genuinely not want to emphasize tennis facilities, on the grounds that she would probably have no one to play with anyway.
2. *Encouraging assimilation to facilitate the reaching of agreement.* A different reason why  $m_1$  may assimilate her preferences to those of  $m_2$  is simply a desire to minimize conflicts that may make it more difficult for the group to find a solution. This pattern is especially likely in cases where  $m_1$  was originally more or less indifferent between two possible preference specifications, before seeing that  $m_2$  has chosen the other one of them. The difference between this case and the previous one is that here,  $m_1$ ’s true preference has not changed, but she has strategically changed her specification of it.

In a similar vein, the more recently developed vacation recommender system CATS (McCarthy et al. [28]; McCarthy et al. [29]) allows group members to achieve some awareness of each other’s activities as they explore vacation options, working simultaneously around a DIAMONDTOUCH table, an environment that facilitates synchronous work of group members on a common project (cf. Dietz and Leigh [12]). In the



**Fig. 20.2.** Illustration of several ways in which the CATS system enhances mutual awareness among group members as they plan a skiing vacation. (Explanation in text. Figure 1 of McCarthy et al. [28], reproduced with permission.)

overview in Figure 20.2, several examples can be seen: Each mountain icon in the large map (a) represents one of the resorts that is being considered, the size of the icon reflecting the currently estimated overall preference of the group for that particular resort. In the description of a particular hotel (b) the check mark or question mark next to the color-coded icon for each group member indicates that group member's estimated interest in the hotel. The color-coded snowflakes on the map (c) indicate what resort each member is investigating at the moment. Finally, each member can send a "critique" that he or she is working on to the other members, thereby sharing his or her thoughts about a particular option.

Although I-SPY's preference specification is largely implicit, there are some phenomena involved in the use of I-SPY that are similar to those that arise with collaborative preference specification of the type we have seen with the TRAVEL DECISION FORUM. These are related to the fact that each user sees the effects of the choices made by other users, even if he does not recognize these effects as such. I-SPY exposes the learned preferences of its community to searchers, in part by highlighting *promoted* results in a search result list (see Figure 20.3 for examples and Smyth et al. [36], and Smyth et al. [37] for further details). Thus, just as in the TRAVEL DECISION FORUM a user can intentionally copy the preferences specified by another group member, in I-SPY, the choices (and thus the implicit preference specifications) of each community member will tend to be affected by the choices of previous searchers. In recent versions of I-SPY, the current user can even see which individual users were responsible for the promotion of a particular link (see Section 20.4.2). The purpose of the provision of this type of information is to provide a sort of explanation of the recommendation of a given



the smarter way to search

Communities About I-Spy

ijcai Search PRIVATE SEARCH

changing worlds: Your Search for ijcai returned 37 Results | Displaying 1 - 20 | Result Page: 1 2 Next

Related Information Search Results for the changing worlds community

**Recent Queries** [VIEW ALL]

1. djuice
2. using cron
3. unix wc command
4. number of occurrences of...
5. wsd/ interface

**Recent Web Pages** [VIEW ALL]

1. main
2. Newbie: Intro to cron
3. perfaq5 - Files and Fo...
4. Guide to Linux
5. The UNIX Guide

**Popular Queries** [VIEW ALL]

1. nragongdo
2. cibenix
3. changingworlds o2
4. dublin map
5. ijcai

**Popular Web Pages** [VIEW ALL]

1. Welcome to UCD
2. Mozilla - Home of the F...
3. Java 2 Platform SE v1.4.2
4. Douglas Newman Good
5. XE.com - The Universal...

**IJCAI-05**

text-only. Nineteenth International Joint Conference on Artificial Intelligence. EDINBURGH, SCOTLAND 30 JULY - 5 AUGUST 2005 ... EDINBURGH, SCOTLAND 30 JULY - 5 AUGUST 2005. accepted papers | poster | tutorials ... Call for papers: closed. 2005 IJCAI awards. Invitation to exhibit ...  
<http://ijcai05.csd.abdn.ac.uk>

**Related Queries** [ijcai 2005 ai conference](#)

**Welcome to International Joint Conference on Artificial Intelligence - 2007**

Registration. Calls. Venue. Feedback. About Us. IJCAI is the International Joint Conference on Artificial Intelligence, the main international gathering of researchers in AI. ... IJCAI and the national AI societie(s) of the host nation(s). The twentieth International Joint Conference on Artificial Intelligence will be held in Hyderabad, India in August 2007 ...  
<http://www.ijcai-07.org>

**Related Queries** [ijcai 2007](#)

**Support Vector Machine Workshop at IJCAI '99**

Workshop on Support Vector Machines at the. International Joint Conference on Artificial Intelligence (IJCAI-99), 2 August 1999, Stockholm, Sweden. NEW! - Preliminary Schedule. 08.30 - 08.35. Welcome (Alex Gammerman) ... pages, and must be submitted in the IJCAI format. Papers must be in a ...  
[http://www.dcs.rhnc.ac.uk/research/compint/areas/comp\\_learn/sv/ijcai.html](http://www.dcs.rhnc.ac.uk/research/compint/areas/comp_learn/sv/ijcai.html)

**Related Queries** [support vector machine workshop ijcai 1999](#)

**IJCAI-99 Home page**

IJCAI-99 Program. Program. (PDF version) Registration information. Related events. Other information. Organizers. Travel guide. City Conference Center, Stockholm, Sweden. July 31 - August 6, 1999 ... Updated

**Fig. 20.3.** Screen shot of I-SPY illustrating several ways in which the current user is helped with information derived from the behavior of other community members. (The “eyes” icon indicates that a result has been promoted to a higher position in the result list than it would have had otherwise.)

link; this function will be discussed in more detail in Section 20.4; but this property is relevant here in that it can lead to (a) a user thinking twice about whether to choose a particular link because of knowing that others will see that he or she chose that link; and (b) the tendency of users to be influenced by the fact that particular other users have already expressed a degree of interest in a given link.

The idea of collaborative preference specification could be extended to some systems that do not yet to make use of it. For example, suppose that  $m_1$  and  $m_2$  jointly make use of MOVIELENS’s *buddy* feature, which is the more recent implementation of the ideas introduced with POLYLENS (cf. <http://movielens.umn.edu/>). It is likely that the movie tastes of  $m_1$  and  $m_2$  overlap to a greater extent than the interests of an arbitrary pair of movie-goers. How might  $m_1$  benefit from being able to use  $m_2$ ’s ratings as a starting point for her own ratings? The most obvious case would be the one in which many of  $m_1$ ’s ratings coincide exactly with those of  $m_2$ ; but even simply knowing which movies  $m_2$  has rated at all might be helpful: Perhaps the most tedious thing about using MOVIELENS is the job of finding movies that you have seen and so are able to rate—a process that may require scrolling through a list of movies that extends over many web pages. A list containing the set of movies that  $m_2$  has rated but  $m_1$  has not rated might contain a higher proportion of movies that  $m_1$  will be able to rate.

## 20.3 Generating Recommendations

No matter how a group recommender acquires the members' preferences, the recommendation for the group will in general be based on information about the preferences of the individual group members. Therefore, some type of *aggregation* method is required, by which information about individual preferences is combined in such a way that the system can assess the suitability of particular items for a group as a whole.<sup>3</sup> The need to choose an aggregation method is the most obvious and intensively studied difference between group recommendation and recommendation for individuals. The topic of preference aggregation is a multifaceted and complex one that has been addressed in various scientific fields (see, e.g., Arrow [2], for a seminal contribution and Masthoff [24] for a summary of this literature from the perspective of group recommendation). The ideas to be discussed in this section overlap to some extent with ideas from these other fields, but a number of new elements are introduced by the technical and practical context of interactive adaptive systems. For example, much of the literature on preference aggregation considers large communities, whose members mostly do not know each other (e.g., citizens of a country who are voting in an election); and the practical contexts require aggregation methods that are technically simpler than many of those that are feasible in adaptive web-based systems.

In this section, we will first discuss ways in which the aggregation problem has been handled in group recommender systems to date. We will then discuss various complications that have not yet received as much attention as they will ultimately require.

We assume for now that the system is supposed to make recommendations concerning just one decision (e.g., one film that is to be watched by the group). The case where a sequence of decisions is to be made (e.g., concerning several TV programs that will be watched on a given evening) will be considered in Section 20.3.3.

### 20.3.1 Approaches to Preference Aggregation

Although the various approaches differ in the ways in which they gather and represent users' preferences, almost all approaches make use of one of three schemas: (a) merging of sets of recommendations, (b) aggregation of individuals' ratings for particular items, or (c) construction of group preference models.<sup>4</sup>

**Merging of Recommendations Made for Individuals.** In cases where what is to be presented is a *set* of candidate solutions, among which the group is to select one for

<sup>3</sup> It is not actually logically necessary for group recommendations to be based on information about the preferences of individual members. For example, a group movie recommender might somehow acquire the knowledge that Walt Disney movies tend to be suitable when parents are taking their small children to a movie, even if the recommender system has no information about how parents and children, respectively, tend to evaluate Walt Disney movies. But since in practice almost all group recommenders start with information about preferences of individual members, we will view the problem of making recommendations for a group as involving preference aggregation.

<sup>4</sup> Somewhat similar distinctions among aggregation approaches have been made by, among others, O'Connor et al. [30]; Kay and Niu [22]; and Yu et al. [39].

TITLE	GENRE	REVIEWS	GROUP	YOUR	cosley@cs.uwm.edu	cosley@quasar
<a href="#">Pixote (1981)</a>	Drama		★★★★★	★★★★★	★★★★★	★★★★★
<a href="#">Wrong Trousers, The (1993)</a>	Animation, Comedy		★★★★★	★★★★★	★★★★★	★★★★★
<a href="#">After Life (1998)</a>	Drama		★★★★☆	★★★★☆	★★★★☆	★★★★★
<a href="#">King of Masks, The (Bian Lian) (1996)</a>	Drama		★★★★☆	★★★★★	★★★★☆	★★★★★

**Fig. 20.4.** Example of a display of group recommendations in POLYLENS. (Adapted with permission from an image supplied by John Riedl. Explanation in text.)

adoption, a simple aggregation method is that of generating a small number of recommended solutions for each member and then merging them into a single list:

1. For each member  $m_j$ :
  - For each candidate  $c_i$ , predict the rating  $r_{ij}$  of  $c_i$  by  $m_j$ .
  - Select the set of candidates  $C_j$  with the highest predicted ratings  $r_{ij}$  for  $m_j$ .
2. Recommend  $\bigcup_j C_j$ , the union of the set of candidates with the highest predicted ratings for each member.

This method was one of those considered for the POLYLENS system (O'Connor et al. [30]). Because of the simple relationship to the generation of recommendations for individuals, this method can be implemented easily as an extension of a recommender for individuals, making use, for example, of any explanation facilities of the individual recommender. In particular, if each recommendation is accompanied by a display of the (predicted) ratings of each member, the members may have quite a good basis for choosing a truly acceptable solution. On the other hand, this approach does presuppose that the members will play an important role in the final decision making, since the list of recommendations does not in itself indicate which solutions are best for the group as a whole. In fact, in the worst case each proposed solution might be excellent for one member but terrible for all of the others. More generally, this method ignores the set of solutions that are not expected to be especially appealing to any member but which might represent the best solution for the group as a whole.

Since this method is rarely even considered for use in group recommenders, we will not discuss it further.<sup>5</sup>

**Aggregation of Ratings for Individuals.** This commonly applied approach starts with the assumption that, for each candidate item  $c_i$  and each member  $m_j$ , the system can predict how  $m_j$  would evaluate (or *rate*)  $c_j$  if he or she were familiar with it:

1. For each candidate  $c_i$ :
  - For each member  $m_j$  predict the rating  $r_{ij}$  of  $c_i$  by  $m_j$ .
  - Compute an aggregate rating  $R_i$  from the set  $\{r_{ij}\}$ .
2. Recommend the set of candidates with the highest predicted ratings  $R_i$ .

<sup>5</sup> A closely related method is considered by Yu et al. [39] for the recommendation of *sequences* of TV programs; essentially the same drawbacks apply in that context as well.

This approach is illustrated by POLYLENS, as can be seen in Figure 20.4. The three right-hand columns in the screen shot display ratings that have been predicted for individual users via the same collaborative filtering method that MOVIELENS uses for individual users. The column labeled “GROUP” shows the aggregated rating. In POLYLENS, the aggregation method is very simple:

$$R_i = \min_j r_{ij}. \quad (20.1)$$

That is, instead of looking for the movie with the highest average rating, POLYLENS applies the strategy of “least misery”: It bases its recommendation on the lowest predicted rating for each candidate, preferring candidates for which the lowest predicted rating for any group member is relatively high. Other plausible aggregation methods will be mentioned in Section 20.3.2.

**Construction of Group Preference Models.** The second widely applied approach to aggregation does not involve any predictions of ratings of individual users. Instead, the system somehow uses its information about the preferences of individual group members to arrive at a model of the preferences of the group as a whole:

1. Construct a preference model  $M$  that represents the preferences of the group as a whole.
2. For each candidate  $c_i$ , use  $M$  to predict the rating  $R_i$  for the group as a whole.
3. Recommend the set of candidates with the highest predicted ratings  $R_i$ .

With regard to Step 1: There are even more possible methods for the construction of group preference models than for the aggregation of individual ratings, since group preference models can take many different forms.

In some cases, the group preference model can be seen as an aggregation of individual preference models. An example is given by LET’s BROWSE: Each individual user’s profile is a set of keyword/weight pairs that reflects the typical content of the pages that this user likes to view. The system computes a model of the group by forming a linear combination of these individual models. From then on, it no longer has to consult the individual models when making recommendations. Similarly, in the context of an in-vehicle multimedia recommender (Yu et al. [40]) and a TV program recommender (Yu et al. [39]), Yu and colleagues introduced and evaluated a method for constructing a group preference model on the basis of individual preference models, using a notion of *distance* between preference models.

In other cases, a group preference model may represent an aggregate of preference models for subgroups, rather than for individual members. This approach is exemplified by the system INTRIGUE (Ardissono et al. [1]), which is designed to help tour guides who need to design tours for heterogeneous groups of tourists that include relatively homogeneous subgroups (e.g., “children”). The tourist group leader divides the tour group into several categories of homogeneous users and specifies a preference model for each such subgroup. The group model is then a weighted average of the subgroup models, with the weights reflecting the importance of the subgroups (e.g., the subgroup of disabled persons is considered especially important because of the special requirements of its members).

The TRAVEL DECISION FORUM takes the focus on group models one step further: In fact, the main function of the system is to help the group members, for each aspect of the vacation that the group members are planning, to arrive at a group preference model that all members have agreed to—that is, at a way of filling out each preference specification form (such as the one shown in Figure 20.1) in such a way that it reflects the preferences of the group as a whole. If we look at the system in this way, the system can be seen as one that recommends specific preferences for the group model (e.g., a rating of ++ for the attribute “Sauna” in Figure 20.1).

I-SPY creates a group (or community) preference model directly on the basis of data concerning the behavior of individual group members, bypassing the level of individual preference models—partly because of privacy considerations, as will be discussed shortly. I-SPY’s basic community preference model consists of a record of queries that have been submitted (by the community of searchers) and the result pages that have been selected for these queries, along with frequency information for these selections. When deciding to what extent to promote a particular search result for a particular community, I-SPY bases its decision on an estimate of how relevant this result page is likely to be for the current query. This estimate is based on the frequency with which this page has been previously selected by community members for the current query and for similar queries.

**Choosing Between Rating Aggregation and Group Preference Models.** Constructing a preference model for the group has the clearest advantages when the group members will have an opportunity to examine and/or negotiate about the group’s model before or after it is actually applied. In this case, for example, the users of INTRIGUE or the TRAVEL DECISION FORUM could settle among themselves once and for all the relative priorities of historical interest and entertainment, instead of debating this issue with respect to each individual attraction. This type of process will be discussed further in Sections 20.4 and 20.5.

If, on the other hand, the group model will be created and applied in the background, without inspection by the group members, the question of whether a group model is better is a more technical one that involves considerations such as efficiency and the quality of recommendations. For example, O’Connor et al. [30] discuss various ways in which POLYLENS could have been designed to create a model of each group (e.g., a “pseudo-user” who represents the interests of the group as a whole) before any recommendations were generated—and some typical consequences of such group models. For instance, a group model might (accurately or not) recommend a movie for which the predicted rating of each individual member was low—something that cannot happen with recommendation-level aggregation.

Another advantage of a group preference model concerns its potential privacy benefits. Recording and maintaining individual user profiles will typically raise privacy concerns, especially if these profiles are owned by some third-party system on the server side. In contrast, the use of a group preference model may go a considerable way toward alleviating these privacy concerns. I-SPY is a case in point. Our web search behavior can be surprisingly revealing when it comes to understanding the likes and dislikes of an individual—far more revealing and valuable to an eavesdropper than movie or music preferences, for example. I-SPY’s use of a community-based profile, in which the

search behavior of individual searchers is merged, means that the search preferences of any individual searcher can no longer be reconstructed.

### 20.3.2 Alternative Goals and Procedures for Aggregation

Even once a general approach has been chosen, the question arises of what particular computational procedure (or *mechanism*) should be used for the aggregation. This is the single question in this area that has received the most attention. The problem is that there are a number of goals that may be desirable in any given situation (e.g., total satisfaction, fairness, and comprehensibility), and conflicts between them can easily arise. In this section, we give several examples of such goals.

Whereas many treatments of these issues (see, e.g., Masthoff [24]; Yu et al. [39]) devote considerable attention to mathematical formulas, quantitative examples, and technical concepts, we will focus on the basic underlying issues and concepts and how they relate to realistic application scenarios.

**Maximizing Average Satisfaction.** Suppose at first that we are taking the approach of aggregating individual ratings. In this case, the goal of maximizing average satisfaction can be achieved by an aggregation function that computes some sort of average of the predicted satisfaction of each member for use as a basis for the selection of candidates (see Equation 20.2). The POCKET RESTAURANTFINDER (McCarthy [26]; cf. Section 20.2.1) applies a variant of this formula to the predicted ratings of restaurants by members of a group who are preparing to go out to dine together. The G.A.I.N. system of Pizzutilo et al. [32], which presents news items on a wall display or an information kiosk, uses a more complex variant of this formula that takes into account uncertainty about which users will be viewing the display at any given time; a similar procedure is applied in FIT (Goren-Bar and Glinansky [14]), which recommends TV shows for members of a family.

$$R_i = \text{average}(\{r_{ij}\}) = 1/n \cdot \sum_{j=1}^n r_{ij}. \quad (20.2)$$

If the predicted ratings are not thought to represent satisfaction accurately, some transformation of them can be used, such as the square of the rating; some results concerning transformations of this sort are given by Masthoff [24].

**Minimizing Misery.** Even if the average satisfaction is high, a solution that leaves one or more members very dissatisfied is likely to be considered undesirable. Even the most ego-centered group member may not want to have to interact with another member who is thoroughly dissatisfied; and such a member may refuse to go along with the solution in any case. In POLYLENS, the minimization of misery is the only criterion applied (see Equation 20.1 above). It is also possible to take this factor into account as a constraint that must be fulfilled by a solution: The lowest predicted rating must not fall below a given threshold.

**Ensuring Some Degree of Fairness** For similar reasons, a solution that satisfies everyone just about equally well is in general preferred to one that satisfies some at the expense of others—all other things being equal. Even more than in the case of minimizing misery, the goal of ensuring fairness is in general combined with some other goal. After all, no-one wants a perfectly fair solution that makes everyone equally miserable. For example (again assuming the approach of aggregating individuals' ratings), the aggregation of predicted individual ratings might include a penalty term that reflects the amount of variation among the predicted ratings, as in Equation 20.3:

$$R_i = \text{average}(\{r_{ij}\}) - w \cdot \text{standard-deviation}(\{r_{ij}\}), \quad (20.3)$$

where  $w$  is a weight that reflects the relative importance of fairness.

**Treating Group Members Differently Where Appropriate.** In some situations, it is generally agreed that the preferences of some group members need to be treated differently than those of others. If two hosts are planning a visit to a restaurant with a visitor from out of town, they are likely to give high priority to the visitor's preferences, requiring only that the solution is not entirely unsatisfactory for themselves (cf. e.g., Kay and Niu [22]). In INTRIGUE, the tourist guide is able to assign higher weights to subgroups such as those of disabled persons or children, on the assumption that these group members are less able to put up with solutions that are even partly unsatisfactory for them.

**Discouraging Manipulation of the Recommendation Mechanism.** The problem of *manipulation* is illustrated by experience with an early version of the system MUSICFX (McCarthy and Anagnost [27]), one of the earliest group recommender systems, which automatically selected music genres for the music to be played in a fitness studio: Although the system essentially applied an averaging procedure to construct a group model from individual preference models, the system also enforced a constraint of the type mentioned above in connection with the "least misery" criterion: Any music genre that was "hated" by any member currently in the gym was removed from the list of possible genres to play. Some users were observed to force an immediate change of genre by adapting their specifications to indicate that they "hated" the genre currently being played—even if they really didn't mind it but simply liked it less well than some other genres.

The potential for manipulation is even more obvious in the TRAVEL DECISION FORUM, in which one group member can often see the preferences specified by the other members. For example, suppose that in Figure 20.1 Claudia's true preference regarding the presence of a sauna was  $\sim$  ("Don't care"): Instead of selecting the middle box in the scale, she might be inclined to select the left-most box (indicating strong disapproval of the availability of a sauna), so as to compensate for the positive preferences specified by the other group members Ritchie and Tina, expecting that the aggregated group preference for a sauna will end up being closer to her own.

When this type of insincere specification of preferences occurs, the aggregation algorithm used will be operating on false premises, since the algorithms presuppose that a group member's expressed preferences reflect his or her true preferences.

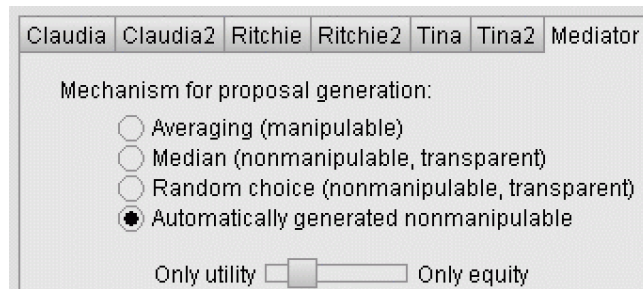
One way of making manipulation difficult is to make it impossible for users to see each others' preferences before specifying their own: If you don't know what the other members prefer, it is hard to distort the resulting recommendation in your own direction by specifying an insincere preference. But users may be able to guess other members' preferences (at least roughly); and in any case, as we have seen, there are advantages to allowing members to see each others' preferences at an early stage.

Manipulation is most likely to be possible if the input that the system uses for making its predictions consists of explicit preference specifications; with implicit inference of preferences, users are much less likely to be able to see how they could influence a recommendation by acting in some particular way. But exceptions can occur; for example, in I-SPY, user can quickly notice that, when they choose a particular link for a given query, that link gets promoted in the search result list for that query. It is then an obvious next move to click on links that one would like to promote (e.g., pages written by the user), regardless of their actual relevance to the query. One can view this type of manipulation as an alternative form of search engine spam, because subsequent users will see potentially irrelevant results being unjustifiably promoted to positions of prominence. As a potential solution to this problem Briggs and Smyth [3] propose the use of an explicit model of trust that provides a filtering mechanism with a view to eliminating the contributions of these manipulative selections: The selections of individual users are evaluated for their reliability. In the simplest sense, a result selection is considered to be reliable if the same link is subsequently reselected by a certain minimum number of searchers for similar queries in the future. This information is used for (among other things) the evaluation of the trustworthiness of individual users, so that recommendations that stem from the activities of users with low trust values can be eliminated or demoted. Preliminary evaluation results suggest that the technique is capable of improving recommendation accuracy.

A different approach to discouraging manipulation is to have the system use an aggregation method that is inherently *nonmanipulable*: It is never in the interest of a given user to specify any preference other than the one that he or she really has. To return to the example with the TRAVEL DECISION FORUM given above: A simple nonmanipulable aggregation mechanism uses as a preference for the group as a whole concerning a given attribute the *median* of the individual preferences for that attribute (i.e., the one that falls exactly in the middle of an ordered list of all preferences). In our example, Claudia will not be able to drag down the group preference for a sauna below + by specifying a low preference herself, since the median preference will be + for any preference that she specifies between -- and +. (It will be left as an exercise for the reader to verify that, with the use of the median mechanism in this setting, no group member could ever benefit by specifying a preference insincerely.)

In general, many nonmanipulable mechanisms may exist for any given preference aggregation problem. In the research area of *automated mechanism design* (see, e.g., Conitzer and Sandholm [8]; Conitzer and Sandholm [9]; Sandholm [33]), methods are developed for automatically generating aggregation methods for a particular setting that (a) satisfy the constraint of being nonmanipulable (at least in that particular setting) and (b) also respect other constraints as well (e.g., maximizing average satisfaction and/or ensuring a certain degree of fairness). The methods introduced by Conitzer and Sand-





**Fig. 20.5.** Part of a dialog box in the TRAVEL DECISION FORUM via which an aggregation procedure can be selected. (The slider, which is active only when a mechanism is to be generated automatically, determines the relative weight of average satisfaction and fairness—cf. Equation 20.3. Adapted with permission from Figure 2 of Jameson [18].)

holm were implemented in the TRAVEL DECISION FORUM, along with the median mechanism just mentioned and two other mechanisms. The system administrator or an individual group member can request that an automatically designed mechanism be used for the generation of recommendations and can specify the properties that this mechanism should have (see Figure 20.5); the system then generates on the fly a mechanism that fulfills the specified constraints. The main issues that arose concerning the appropriateness of such automatically designed mechanisms were whether they were sufficiently comprehensible (or “transparent”, in the terms of Figure 20.5) and acceptable to users. These issues will be discussed in the next subsection.

**Ensuring Comprehensibility and Acceptability.** As will be discussed in Section 20.4, group members sometimes like to be able to understand the rationale behind a recommendation. In particular, they may want to check to what extent acceptability criteria such as the ones discussed earlier in this section are being fulfilled. Even with ingenious visualizations such as those that will be shown in Section 20.4, it may be difficult for a system to explain a recommendation if the mechanism by which the recommendation was generated is inherently complex and/or counterintuitive. Therefore, it may be worthwhile to choose an inherently comprehensible mechanism even if it is not the best mechanism in terms of the other criteria.

An example of a comparison of mechanisms in terms of their inherent comprehensibility and acceptability is the exploration of automatically designed nonmanipulable mechanisms in the context of the TRAVEL DECISION FORUM (cf. the previous subsection and Jameson et al. [20]). One fundamental limitation of the automatically designed mechanisms is the fact that such a mechanism cannot be represented with a simple formula (such as the formula for the average or the median) but rather has to be represented by a table that specifies, for each possible combination of preferences of individual users, which item should be chosen.<sup>6</sup> Therefore, a group member cannot apply the mechanism mentally in order to predict or understand recommendations.

<sup>6</sup> Actually, automatically generated mechanisms are often nondeterministic: For each possible combination of preferences, the mechanism specifies a vector of probabilities associated with

Moreover, unless special acceptability constraints are applied in the mechanism generation process, the generated mechanism may give rise to recommendations that strike people as counterintuitive and inappropriate (e.g., proposing for some combinations of individual preferences an outcome that none of the group members likes, even though there exist outcomes that some members like). In sum, automated mechanism design is an approach that deserves further attention, but special attention must be paid to the goal of ensuring adequate comprehensibility and acceptability.

### 20.3.3 Further Complications Concerning Preference Aggregation

The often conflicting goals discussed in the previous subsection are in themselves enough to make the problem of choosing a suitable aggregation procedure a difficult one. But there are additional complications that need to be taken into account in some settings.

**Generating Recommendations Concerning Multiple Decisions.** So far, we have been focusing on the situation in which a recommender will make recommendations to a group concerning just one decision. But often the group members will expect a system to make recommendations concerning a larger set of decisions, either at the same time or in succession: INTRIGUE's tour guide will choose several sights to visit; the music selection systems will choose one song after the other; and a TV program recommender will recommend several programs for a group to watch in succession.

In this type of situation, the procedure for generating recommendations about the entire sequence can be related to a procedure with respect to individual decisions in any of several ways. Figure 20.3 compares three approaches, which differ in terms of several criteria.

1. The simplest approach is to treat each decision separately, ignoring the fact that there will be a sequence of decisions. Because of its simplicity, this approach tends to be computationally simple and easy to explain to users. One drawback is that a goal such as ensuring fairness can be taken into account only with respect to individual decisions, whereas it can be advantageous to consider it with regard to the entire set of decisions. For example, it may seem fair enough to recommend a single TV program that is much less attractive for one group member than for the others as long as that group member's overall satisfaction with the sequence of programs is comparable to that of the other members. Trying to ensure approximately equal satisfaction among a group members with each individual program in the sequence may rule out too many options that would be attractive for the group as a whole. An even clearer limitation of this approach arises in cases where the system can acquire feedback about the results of each decision before making recommendations concerning the next one. Suppose, for example, that one group member was especially dissatisfied with a given TV program, even though this dissatisfaction was not predicted in advance. It may be feasible and desirable to bias the recommendations concerning one or more subsequent programs in favor of that group member, so that he or she ultimately reaches an appropriate level of satisfaction; but

---

the various possible outcomes, which is to be used for the random selection of one of the outcomes.

**Table 20.3.** Positive (+), negative (-), and intermediate (+/-) aspects of three approaches to the treatment of a sequence of decisions by a group recommender.

Criterion	<i>Approach to treating a sequence of decisions</i>		
	Independently	As one complex decision	Individually but with consideration of other decisions
Computational complexity	+	-	+/-
Comprehensibility	+	-	+/-
Appropriateness of evaluation criteria	-	+	+/-
Ability to take into account actual results of individual decisions	-	-	+
Applicability when decisions and members involved are not known in advance	+	-	+/-
Ability to take into account additional ordering constraints	-	+	+/-

this type of compensation is not possible if the decisions are treated completely independently. Similarly, this approach cannot deal with other constraints that concern relationships among members of the sequence (e.g., the possible undesirability of presenting two very similar items in succession, cf. the discussion of FLYTRAP below; or a possible tendency of earlier items in the sequence to have a generally larger impact on users than later items, cf. Masthoff [25]).

2. The opposite approach is to view the recommendations for the entire sequence as a single recommendation problem, much like that of recommending complex items (such as vacations) that differ with respect to a number of attributes. This approach is applicable only if it is known in advance what particular sequence of decisions is going to be made and what group members are going to be involved (e.g., how many times and on which particular occasions a group of diners is going to go out to dine together). This approach makes it straightforward to apply criteria such as fairness to entire sequences. On the other hand, it does not make it possible to take into account the results of previous decisions, since the decision making process for the entire sequence is completed before any decisions are executed. Also, the necessary computational procedures tend to be more complex; for example, the number of possible sequences may be too large for it to be feasible for the system to iterate through all possible sequences and evaluate their suitability.

3. An approach that lies between these two extremes starts with the idea of treating each decision problem separately but makes some adjustments to take into account the

fact that a sequence is being dealt with. For example, each individual decision might be approached with the goal of maximizing overall satisfaction; but if the system notices that, up to a given point in time, a given member has been less satisfied than the others, his or her satisfaction can be given greater weight in subsequent decisions, until the discrepancy has been eliminated. This approach is able to take into account the actual results of decisions (as opposed to only the predicted results). Although more complex than independent treatment of decisions, the method may be reasonably explainable if it corresponds with familiar decision making schemas from everyday life (e.g., “That last program was awful for Mary, so let’s give her a break this time.”) This approach does not require the set of decisions to be known in advance, but it can be applied most effectively if a good deal is known. For example, if the system has decided to grant a certain amount of extra satisfaction to a particular group member while making the subsequent recommendations, it will be helpful to know how many decisions remain to be made.

An example of this third approach is found in the FLYTRAP music selection system (cf. Section 20.2.1): The system has to choose songs one at a time, because the set of persons who are present to hear them frequently changes. But its selection procedure does take into account constraints imposed by a “DJ agent” that tries, for example, to avoid abrupt and distracting changes of genre.

#### 20.3.4 Preference Specifications That Reflect More Than Personal Taste

In most analyses, it is assumed that the preferences specified by a group member represent simply the desires of that individual member (e.g., in a TV context, the programs that the group member would watch if he or she were watching alone). But in some settings, a group member may be taking into account the assumed interests of other members when expressing his or her own preferences. For example, we noted in connection with the TRAVEL DECISION FORUM that users sometimes expressed preferences in such a way as to minimize the likelihood of conflict. To take a more extreme example, consider a mother who is taking her two children to the movies and whose primary motivation is to find a movie that the children will like and that she herself will not hate having to sit through. In this case, her preference specification (e.g., a high rating for a particular children’s movie) will reflect only to a minimal extent her own taste in movies. So it would not be appropriate, for example, for a recommender system to look straightforwardly for a “compromise” between the mother’s expressed preference and the preferences expressed by the children (though some more sophisticated aggregation of the various expressed preferences might well make sense).

An additional complication arises when the group members’ preferences reflect not just subjective tastes but also knowledge that may be relevant to the choices of the other members as well. For example, suppose that a member  $m_1$  of a travel group who is especially familiar with Switzerland expresses a strong preference for a given Swiss ski resort: The other group members may be willing to give extra weight to  $m_1$ ’s opinion, even if their own evaluation criteria for resorts are somewhat different, on the grounds that the resort in question is especially likely to be good at least according to  $m_1$ ’s

criteria, which overlap to some extent with the criteria of the other group members.<sup>7</sup> An ad hoc way of taking differences in knowledge into account is to assign greater weight to the preferences of more knowledgeable group members; but this method does not address in a principled way the relationship between knowledge and preferences.

### 20.3.5 At What Points Can These Complexities Be Dealt With?

After this lengthy discussion of conflicting goals for preference aggregation procedures and additional complexities, the reader may be wondering whether it will ever be possible to deal with these issues adequately in group recommender system. Fortunately, the issues can be dealt with by different people at different points in time:

1. *The system's designers and/or deployers may specify an appropriate means of handling each problem:* The persons who are designing a group recommender—or arranging the deployment of an existing recommender in a given context—can consider each of the issues discussed above with regard to their particular target group and application setting and work out some locally appropriate solution. For example, the designers of POLYLENS thought that the “least misery” aggregation function would be appropriate because they expected most groups of people who go to see a movie together to be small (i.e., 2 or 3 members); for settings involving larger groups, the same function would probably lead to too many cases in which a solution that would be liked by many members would in effect be vetoed by the one person who liked it least. If the designers of a restaurant recommender anticipate that there will often be some group members who are familiar with the restaurants in question and others who are not, they might look for a principled way of treating the two types of group member differently.
2. *The system's users may select a suitable preference aggregation method for each decision:* A system can allow the users to decide what aggregation mechanism is to be used, either before any recommendations are made or during an iterative process of requesting recommendations and adjusting the aggregation function. For example, with INTRIGUE the tour guide can specify a different set of subgroup weights for each tour group. As was mentioned above, a variety of aggregation mechanisms can be chosen in the TRAVEL DECISION FORUM. This idea could be adopted in many recommendation settings.
3. *Users can take any remaining factors into account when evaluating specific recommendations and negotiating about the final decision:* If the system presents a number of recommendations and allows users to choose which one(s) they want to adopt, it may not be necessary for the system itself to deal with all of the subtle problems that can arise. Instead, the users themselves may be able to take these issues into consideration, making use of their long experience with social interaction and relationships. After all, even the most subtle of the issues discussed in this section concern matters that people are accustomed to dealing with in everyday life. If

---

<sup>7</sup> Hastie and Kameda [16] show how aggregation procedures can be compared in terms of their suitability for arriving at accurate decisions even in settings where the group members are assumed to have identical preferences—that is, where knowledge aggregation rather than preference aggregation is involved.

the group recommender is designed on the assumption that there are some aspects of the decision problem that are better dealt with by the users themselves, the function that the system serves is mainly decision support rather than decision making. In these cases, special importance should be assigned to the third and fourth sub-tasks of a group recommender (explaining recommendations and supporting final decision making), which will be discussed in the next two sections, respectively.

## 20.4 Explaining Recommendations

### 20.4.1 Motivation

Given the many ways in which recommendations for a group can be derived—and the often conflicting goals that can be pursued—it is natural that group members should want to understand to some extent how a recommendation was arrived at—and in particular, how attractive a recommended item is likely to be to each individual group member.

Recommender systems for individuals often accompany each recommendation with some sort of analysis of its predicted acceptability; the analysis may range from a simple index of the system's confidence to a complex visualization of the pros and cons of the recommended solution (see, e.g., Herlocker et al. [17], and the chapter in this volume by Schäfer et al. [34]). With group recommenders, it is in principle possible to present such an analysis for each individual member, for the group as a whole, and perhaps for subsets of members. A member  $m_1$  may be interested in the analysis for  $m_2$  because  $m_1$  considers it important that  $m_2$  be satisfied, because  $m_1$  wants to make sure that she is getting “as good a deal” as  $m_2$ , or simply in order to understand how the recommendation was derived.

### 20.4.2 Treatment in Existing Group Recommenders

As can be seen in Figure 20.6, LET'S BROWSE explains each of its web page recommendations by listing the keywords in the page that it assumes to be of interest to all group members. By showing where these keywords are located in each member's profile, the system also allows each user to guess how interesting each member will find the page. In the example in the figure, it looks as if the (hypothetical) group member George Lucas will be less enthusiastic about the page than the member Bill Gates, given that Lucas is only marginally interested in technology. As some of the systems to be discussed below suggest, it might be a worthwhile further step for LET'S BROWSE to present an explicit estimate of the likely interestingness of the page for each group member and perhaps for the group as a whole. In this way, for example, Lucas might more readily accept the system's recommendation of the page involved in Figure 20.6, seeing quickly that it is more interesting to other members than it is to him; or, depending on his overall attitude, he might object to the recommendation for just that reason. On the other hand, since (as was mentioned in Section 20.3.1) LET'S BROWSE uses a group-level model to compute its recommendations, the computation of predictions for individual group

**Let's Browse!**  
collaborative web browsing demo  
Henry Lieberman, Neil W. Van Dyke, and Adriana Vivicqua

This page might interest **Bill, George, and Nicholas** because it concerns **technology** and **travel**.

**Bill Gates**  
Microsoft Corp.  
billg@microsoft.com

**PROFILE BUILT FROM:**  
<http://www.microsoft.com/billgates/>

**PROFILE KEYWORDS:**  
**technology**<sup>(56)</sup> **internet**<sup>(50)</sup> **travel**<sup>(48)</sup>  
windows<sup>(46)</sup> pc<sup>(43)</sup> subsidiary<sup>(39)</sup>  
investment<sup>(32)</sup> ceo<sup>(29)</sup> intellectual<sup>(20)</sup>  
property<sup>(20)</sup> ...

**George Lucas**  
LucasArts Entertainment

**PROFILE BUILT FROM:**  
<http://members.tripod.com/~gnomebasher/lucas.htm>

**PROFILE KEYWORDS:**  
skywalker<sup>(52)</sup> business<sup>(42)</sup> **travel**<sup>(39)</sup>  
force<sup>(30)</sup> star<sup>(25)</sup> wars<sup>(24)</sup> internet<sup>(18)</sup>  
graffiti<sup>(14)</sup> **technology**<sup>(11)</sup> digital<sup>(10)</sup>

**Fig. 20.6.** Screen shot of LET's BROWSE that shows how the system explains a recommendation for three (hypothetical) users of a particular web page. (Adapted with permission from an image supplied by Henry Lieberman.)

members for presentational purposes would involve additional overhead, and it would not reflect the way in which the system actually arrives at its recommendations.

POLYLENS gives an idea of how a display that shows predictions for individuals might look. Since POLYLENS uses collaborative filtering, it cannot explain a movie recommendation in terms of the movie's content; but it does show the predicted rating for each group member and for the group as a whole (see Figure 20.4 above). In addition to explaining each recommendation in terms of the underlying predictions for individuals, this visualization makes it possible for the attentive user to notice how group recommendations are generated—via the “least misery” principle (Equation 20.1)—by comparing predictions for the individual members with the predictions for the group. Incidentally, more than 90% of the users surveyed stated that they had no privacy concerns about having their predicted ratings shown to other group members—a result that encourages the development of additional methods that expose individual-level predictions to all group members.

INTRIGUE offers a type of explanation (Figure 20.7) that is partly similar to that of POLYLENS: It shows the predicted attractiveness of each recommended attraction for the tourist group as a whole; but instead of simply presenting a predicted attractiveness for each subgroup, it explains verbally the aspects of the attraction that are likely to

**INTRIGUE** new technologies for tourist assistance [email us!](#)  
 Dip. di Informatica - Università di Torino - C. Svizzera, 185 - 10149 Torino (Italia)

[back](#) [Advanced Search](#) [View Agenda](#) [Separate listing by groups](#) [Unique listing \(method 1\)](#) [Unique listing \(method 2\)](#)

"Separate listing by group": it shows separate lists, with items sorted on the basis of the different user's preferences  
 "Unique listing (method 1)": it shows a single list, taking into account the needs of the whole group  
 "Unique listing (method 2)": it shows a single list, trying to satisfy a little bit everybody

Suggestions for the whole group:

[add to agenda](#) [Lingotto](#) ★★★ For children it is much eye-catching, it requires low background knowledge, it requires a quite short. For yourself it is much eye-catching and it has high historical value. For impaired it is much eye-catching historical value.

[add to agenda](#) [Palazzo Reale](#) ★★★ For children it is much eye-catching. For yourself it is much eye-catching, it has high high artistic value. For impaired it is much eye-catching and it has high artistic value.

[add to agenda](#) [Palazzo Madama](#) ★★★ For yourself it has high historical value.

[add to agenda](#) [Palazzo Carignano](#) ★★ For children the visit is quite short. For yourself it has high historical value.

[add to agenda](#) [Palazzo Saluzzo di Paesana](#) ★★ For children the visit is quite short.

**Fig. 20.7.** Example of INTRIGUE's main explanation method. (Adapted with permission from Figure 3 of Ardissono et al. [1].)



**Fig. 20.8.** Part of a visualization from the FLYTRAP music selection system. (The songs most likely to be played are shown in the middle of the display. The color coding reflects the estimated degree of interest of the two current listeners Kris and Andy, as well as the influence of the DJ agent. Adapted with permission from Figure 1 of Crossen et al. [11].)

appeal to that subgroup (not mentioning the less appealing aspects). This type of explanation seems helpful for the tour guide who would like all group members to accept the recommendation, but it does not convey a clear idea of how attractive the recommended tour is to each subgroup.<sup>8</sup>

<sup>8</sup> INTRIGUE also offers a different type of explanation that shows only predictions for individual subgroups; see Figure 8 of Ardissono et al. [1].




FLYTRAP offers a completely different type of visualization (see Figure 20.8), which indicates which users are present in the room in which music is being played, what songs are more or less likely to be selected for playing, and how these songs are expected to be evaluated by the users who are present and by the “DJ agent”. Since the users of FLYTRAP do not have an opportunity to influence the choice of songs directly, the purpose of this visualization is to convey a general understanding of how the system works.

I-SPY incorporates a number of strategies for making clear the reasons for the promotion of a given link in a search result list. Since the users of I-SPY are not viewed as working together when they look for relevant links, the focus here is not on showing the desirability of an option for particular group members with a view to resolving conflicts. Nonetheless, it has proven worthwhile to provide information about how other community members have dealt with the page in question in the past, because such information helps the current user to judge its value for him- or herself. Types of information offered include (a) *related queries* for which the page in question has been selected as a promising result (see Figure 20.3); (b) quantitative and temporal information such as “10% of searchers have also selected this result for similar queries as recently as 15 minutes ago” (Coyle and Smyth [10]); and (c) the names of the users who are responsible for the promotion of the page (a by-product of the antimanipulation measures discussed in Section 20.3.2; see Briggs and Smyth [3]).

The TRAVEL DECISION FORUM introduces two novel, complementary methods that aim to provide a more detailed picture of the consequences of a given proposal for each group member:

1. The first method automatically follows from the use of the preference specification form for the presentation of proposals (see Figure 20.1). Since both the specified preferences and the recommended joint preferences are shown on the same set of scales, the user can quickly see which group members should be most / least satisfied with a given proposal (i.e., the ones whose preferences are closest to / farthest from the highlighted cells). Also, with a bit of practice the user can see more complex patterns (e.g., Claudia might notice that “Tina and Ritchie have generally similar preferences, and they usually get their way, while my preferences have little influence”). Any verbal arguments associated with the other members’ stored preferences add further detail to the picture of how they would evaluate a given proposal.

2. The second method takes into account the fact that any graphical explanation of a recommendation is likely to be less interesting, vivid, and memorable than the type of feedback that group members get while they are interacting face to face: A member who is disappointed with a proposal may complain about specific aspects of it in an emotional manner, formulating (or repeating) arguments. In settings where all group members are physically present in front of the group recommender system, this type of face-to-face discussion is likely to occur spontaneously. For settings in which no such direct communication is possible, the TRAVEL DECISION FORUM tries to recapture some of the flavor of face-to-face interaction through animated characters: It is assumed that at any given moment only one group member will be interacting with the system; each of the other members is represented by an animated character who bears that member’s name. Whenever the system (represented by an animated character called



**Fig. 20.9.** Use of animated characters in the TRAVEL DECISION FORUM to represent the likely responses of the absent group members to a proposal. (Here, the representative of Tina is responding to the proposal shown in Figure 20.9. Adapted with permission from Figure 3 of Jameson [18].)

the *mediator*) has recommended a particular joint preference model for a given value dimension (such as “health facilities”), he asks the representatives of the absent group members to comment on it in turn. Parts of a typical performance of a representative are shown in Figure 20.9.

This type of simulated reaction can heighten the group members’ awareness of the other members’ points of view—including their motivational orientations—and overcome the natural tendency to focus on one’s own evaluations. Also, like the explanations of INTRIGUE, these presentations are selective, focusing on the most important considerations for each group member. The user can switch attention back and forth between the animated agents and the graphical explanation, because the two types of explanation make use of largely complementary communication channels.

### 20.4.3 Concluding Remarks About Explanations

These examples from existing systems illustrate (a) that there are many types of information that can be conveyed by an explanation of the system’s recommendations and (b) that the function of such explanations is not necessarily to convince the group members that the system’s recommendations should be accepted. Instead, the system’s explanations are often best seen as information that puts the group members in a better position to make a final decision, which may deviate radically from the recommendations made. The process of arriving at a final consensus is the subject of the next section.

## 20.5 Helping Group Members to Achieve Consensus

Even with a recommender system for individual users, no matter how appropriate and compelling the system’s recommendations and explanations may be, there is usually

no guarantee that any of the recommendations will be adopted. With individual recommenders, although the decision process may be complex, it typically takes place within the mind of a single person. With a group recommender, extensive debate and negotiation may be required, which may be especially problematic if the members are not able to communicate easily.

### 20.5.1 Treatment in Existing Systems

Group recommender systems have tended not to provide explicit support for the process of arriving at a final decision. Such support may in fact not be required if any of the following conditions hold:

1. The system simply translates the most highly rated solution into action without requesting the consent of any users.  
This method is applied by the music selection systems ADAPTIVE RADIO, FLY-TRAP, and MUSICFX, which select and play music autonomously on the basis of the preferences of the group members who are present. It would in fact usually be impractical to allow the persons who happen to be present in a public space to debate about each piece of music that is to be played.
2. It is assumed that one group member is responsible for making the final decision.  
LET'S BROWSE is based largely on the assumption that one group member controls the pointing device and will therefore make most of the decisions about what pages to visit. With INTRIGUE, it appears to be assumed that the tourist guide will decide what tour should be taken.
3. It is assumed that group members will arrive at the final decision through conventional discussion (e.g., face to face or by phone).  
An especially clear example where this assumption is justified is the situation where several group members are working simultaneously with the CATS vacation recommender system (cf. Section 20.2.2) on the DIAMONDTOUCH interactive tabletop. In addition to the usual broad bandwidth of face-to-face communication, the group members can refer during their negotiations to the various information displays provided by the system.

The only other system we know of that specifically supports communication among group members for the purpose of final decision making is the TRAVEL DECISION FORUM—understandably, since this system is designed specifically for groups of users who usually cannot communicate with each other in real time. The animated representatives of absent group members (Section 20.4) do not serve only as a means of visualizing the implications of recommended solutions for the absent members. In addition, each member can grant her representative a certain amount of authority to accept proposals during interactions with another group member. For example, in Figure 20.9, Tina's representative states that she cannot accept the current proposal. If instead the representative had accepted it and the same were true of Ritchie's representative and the current user Claudia, the proposal would have been treated as finally accepted, even though Tina and Ritchie had not really seen it.

### 20.5.2 Possible Extensions to Existing Systems

Since in most applied scenarios the overhead of animated agents would be too great, we should consider how functions such as those of the TRAVEL DECISION FORUM can be realized in a lighter-weight manner. For example, a system like POLYLENS might allow each member to specify that they are willing to go to any movie whose predicted rating for them is above some threshold (e.g., 4 stars out of 5). In that case, the system could present not only recommendations but also a subset of the recommended movies that can be decided on without further consultation; a designated group member could then go on and buy tickets for any of these movies. If it is assumed that each group member will view the recommendations before a decision is made, a procedure could be introduced that allowed the group members to vote for movies among the recommended ones, also indicating which particular ones they are willing to accept. In this case as well, it could be agreed that a designated group member could make the final decision.

As the reader may have noticed, this type of voting mechanism can in itself be viewed as a simple recommender system that makes use of explicit preference specifications and helps the group members to choose among the recommended options. And in fact we could apply all of the concepts introduced so far to this “recommender”, considering, for example, whether the group members should be allowed to see each other’s votes (cf. Section 20.2), how the votes should be counted and weighted (Section 20.3), how the results of the voting should be presented (Section 20.4), and even how the really final decision ought to be made (the present section). Fortunately, we are not faced with an infinite regress, since the recommendation problem that we are dealing with now—choosing from among a small number of recommended items—is considerably simpler than the original problem, and simple solutions may be quite adequate. For example, suppose that for the original recommendations an aggregation function was used which gave especially high weight to particular group members (e.g., the visitors from out of town). It may not be necessary for the final voting mechanism to be biased in their favor as well, since the set of recommendations itself will already contain mainly films that the guests will like; and the hosts are likely in any case to vote in a way that takes into account the greater importance of the guests.

More generally, decisions are often made in several stages, and in each stage a different type of decision making procedure may be applied. Issues that require a great deal of attention in one stage may be simpler to deal with in another stage.

## 20.6 General Conclusions

### 20.6.1 Conclusions Concerning Group Recommender systems

This chapter has shown that the differences between recommending for groups and recommending for individuals are more numerous, important, and complex than one would tend to think at first glance.

1. New methods of acquiring information about users’ preferences are available, especially when group members specify their preferences explicitly; but the interpretation of explicitly specified preferences can be less straightforward than in the case of individual users.

2. The process of selecting items to recommend for the group as a whole can involve much more than the application of numerical formulas, and the appropriateness of the potentially applicable methods depends on various aspects of the application setting.

3. An explanation of the considerations that underlie a recommendation can take many different forms and convey many different types of information, which can be processed further by the group members in their efforts to do justice to considerations that could not be taken into account by the system.

4. The process of arriving at a final decision can require communication and negotiation, which can be supported in various ways by the system, depending on the nature of the setting.

Because of the many specific questions raised by these differences between recommendation for groups and for individuals, it is all too natural for designers and researchers to focus on one or two of the differences, implicitly adopting with respect to other issues a default solution that may be far from optimal for their particular setting. We hope that, by accumulating and organizing ideas and experience concerning the most important differences, this chapter will enable designers and researchers to do justice to all of these differences in their own work, either by adopting ideas that have already been developed or by working out new solutions of their own.

### 20.6.2 Implications for Other Types of Adaptive Web-Based System

Even more generally, just about any type of system that adapts to its users can be seen in some sense as a recommender system, and it may be natural to extend it for adaptation to groups of users. For example, a system for personalized information access (cf. the chapter by Gauch, Speretta, and Micarelli [13]) can be seen as “recommending” particular documents to a user; and it is reasonable to adapt to groups of cooperating, information-seeking users. Either explicitly or implicitly, we will then have to deal with issues such as the specification and aggregation of preferences of the various group members.

Similarly, a system that offers adaptive navigation support (cf. the chapter by Brusilovsky [4]) can be seen as recommending moves within a hyperspace; and it is natural to consider groups of navigating users (as is in fact done in LET’s BROWSE).

With systems that aim to encourage and support collaboration (cf. the chapter by Soller [38]), one of the many functions is that of “recommending” to a set of potential collaborators courses of action that are predicted to be beneficial to them. From a group recommendation perspective, issues that arise include those of (a) how to select a form of collaboration that best takes into account the possibly diverging goals and preferences of the potential collaborators; and (b) how to convince the potential collaborators that they would in fact benefit from the proposed collaboration—or how at least to enable them to devise some form of collaboration that they consider more suitable.

Turning to the more domain-specific topic of systems for health care (cf. the chapter by Cawsey et al. [6]), consider the example of a system designed to persuade members of a family to adopt healthier eating habits: This type of intervention can be seen as involving recommendations to a group of users who have different roles and preferences. Most of the issues discussed in this chapter take a more complex form in this type of setting than in the application settings discussed so far.

In the light of these and many other possible examples, it seems natural that group recommender systems should attract increasing attention not only from designers and researchers who are interested in this particular type of system but also from those who work with other types of adaptive web-based systems.

**Acknowledgments.** For their support for the preparation of this chapter and of their own relevant research, the authors are grateful to the following sources: for Anthony Jameson, the German Ministry of Education and Research (BMBF) (projects MIAU and SPECTER); for Barry Smyth, Science Foundation Ireland under grant 03/IN.3/ I361 and the Informatics Research Initiative of Enterprise Ireland. The anonymous reviewers of this chapter provided valuable feedback that led to substantial improvements.

## References

1. Ardissono, L., Goy, A., Petrone, G., Segnan, M., Torasso, P.: INTRIGUE: Personalized recommendation of tourist attractions for desktop and handset devices. *Applied Artificial Intelligence* 17(8-9) (2003) 687–714
2. Arrow, K.: *Social Choice and Individual Values*. 2nd edn. Wiley, New York (1963)
3. Briggs, P., Smyth, B.: Modeling trust in collaborative web search. In: *Proceedings of the Sixteenth Irish Artificial Intelligence and Cognitive Science Conference*. (2005)
4. Brusilovsky, P.: Adaptive navigation support. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. *Lecture Notes in Computer Science*, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
5. Burke, R.: Hybrid web recommender systems. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. *Lecture Notes in Computer Science*, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
6. Cawsey, A., Grasso, F., Paris, C.: Adaptive information for consumers of healthcare. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. *Lecture Notes in Computer Science*, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
7. Chao, D., Balthrop, J., Forrest, S.: Adaptive Radio: Achieving consensus using negative preferences. In: *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work*. (2005) 120–123
8. Conitzer, V., Sandholm, T.: Complexity of mechanism design. In Darwiche, A., Friedman, N., eds.: *Uncertainty in Artificial Intelligence: Proceedings of the Eighteenth Conference*. Morgan Kaufmann, San Francisco (2002) 103–110
9. Conitzer, V., Sandholm, T.: Applications of automated mechanism design. (2003)
10. Coyle, M., Smyth, B.: Explaining search results. In Kaelbling, L.P., Saffiotti, A., eds.: *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, San Francisco (2005) 1553–1555
11. Crossen, A., Budzik, J., Hammond, K.: Flytrap: Intelligent group music recommendation. In Gil, Y., Leake, D., eds.: *IUI 2002: International Conference on Intelligent User Interfaces*. ACM, New York (2002) 184–185
12. Dietz, P., Leigh, D.: DiamondTouch: A multi-user touch technology. In: *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology*, Orlando, FL (2001) 219–226

13. Gauch, S., Speretta, M., Chandramouli, A., Micarelli, A.: User profiles for personalized information access. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Lecture Notes in Computer Science, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
14. Goren-Bar, D., Glinansky, O.: Family stereotyping – A model to filter TV programs for multiple viewers. In: *Proceedings of the Workshop on Personalization in Future TV at the Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Malaga, Spain (2002)
15. Goy, A., Ardissono, L., Petrone, G.: Personalization in e-commerce applications. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Lecture Notes in Computer Science, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
16. Hastie, R., Kameda, T.: The robust beauty of majority rules in group decisions. *Psychological Review* **112**(2) (2005) 494–508
17. Herlocker, J., Konstan, J., Riedl, J.: Explaining collaborative filtering recommendations. In: *Proceedings of the 2000 Conference on Computer-Supported Cooperative Work*, Philadelphia, PA (2000) 241–250
18. Jameson, A.: More than the sum of its members: Challenges for group recommender systems. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces*, Gallipoli, Italy (2004) 48–54
19. Jameson, A., Baldes, S., Kleinbauer, T.: Two methods for enhancing mutual awareness in a group recommender system. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces*, Gallipoli, Italy (2004) 447–449
20. Jameson, A., Hackl, C., Kleinbauer, T.: Evaluation of automatically designed mechanisms. In: *Proceedings of the First Bayesian Modeling Applications Workshop at the Nineteenth Conference on Uncertainty in Artificial Intelligence*, Acapulco, Mexico (2003)
21. Jennings, N., Faratin, P., Lomuscio, A., Parsons, S., Sierra, C., Wooldridge, M.: Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation* **10**(2) (2001) 199–215
22. Kay, J., Niu, W.: Adapting information delivery to groups of people. In: *Proceedings of the First International Workshop on New Technologies for Personalized Information Access at the Tenth International Conference on User Modeling*, Edinburgh (2005)
23. Lieberman, H., Van Dyke, N., Vivacqua, A.: Let's Browse: A collaborative Web browsing agent. In Maybury, M., ed.: *IUI99: International Conference on Intelligent User Interfaces*. ACM, New York (1999) 65–68
24. Masthoff, J.: Group modeling: Selecting a sequence of television items to suit a group of viewers. *User Modeling and User-Adapted Interaction* **14**(1) (2004) 37–85
25. Masthoff, J.: The pursuit of satisfaction: Affective state in group recommender systems. In Ardissono, L., Brna, P., Mitrovic, A., eds.: *UM2005, User Modeling: Proceedings of the Tenth International Conference*. Springer, Berlin (2005) 296–306
26. McCarthy, J.: Pocket RestaurantFinder: A situated recommender system for groups. In: *Proceedings of the Workshop on Mobile Ad-Hoc Communication at the 2002 ACM Conference on Human Factors in Computer Systems*, Minneapolis (2002)
27. McCarthy, J., Anagnost, T.: MusicFX: An arbiter of group preferences for computer supported collaborative workouts. In: *Proceedings of the 1998 Conference on Computer-Supported Cooperative Work*. (1998) 363–372
28. McCarthy, K., Salamó, M., Coyle, L., McGinty, L., Smyth, B., Nixon, P.: Group recommender systems: A critiquing-based approach. In Paris, C., Sidner, C., eds.: *IUI 2006: International Conference on Intelligent User Interfaces*. ACM, New York (2006) 267–269
29. McCarthy, K., Salamó, M., McGinty, L., Smyth, B.: CATS: A synchronous approach to collaborative group recommendation. In: *Proceedings of the Nineteenth International Florida Artificial Intelligence Research Society Conference*, Melbourne Beach, FL (2006)

30. O'Connor, M., Cosley, D., Konstan, J., Riedl, J.: PolyLens: A recommender system for groups of users. In Prinz, W., Jarke, M., Rogers, Y., Schmidt, K., Wulf, V., eds.: *Proceedings of the Seventh European Conference on Computer-Supported Cooperative Work*. Kluwer, Dordrecht, The Netherlands (2001)
31. Pazzani, J., Billsus, D.: Content-based recommender systems. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Lecture Notes in Computer Science, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
32. Pizzutilo, S., De Carolis, B., Cozzolongo, G., Ambruso, F.: Group modeling in a public space: Methods, techniques and experiences. In: *Proceedings of WSEAS AIC 05, Malta (2005)* 175–180
33. Sandholm, T.: Automated mechanism design: A new application area for search algorithms. In Rossi, F., ed.: *Principles and Practice of Constraint Programming — CP 2003: 9th International Conference*. (2003) 19–36
34. Schafer, B., Frankowski, D., Herlocker, J., Sen, S.: Collaborative filtering recommender systems. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Lecture Notes in Computer Science, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
35. Smyth, B.: Case-base recommendation. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Lecture Notes in Computer Science, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
36. Smyth, B., Balfe, E., Boydell, O., Bradley, K., Briggs, P., Coyle, M., Freyne, J.: A live-user evaluation of collaborative web search. In Kaelbling, L.P., Saffiotti, A., eds.: *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, San Francisco (2005) 1419–1424
37. Smyth, B., Balfe, E., Freyne, J., Briggs, P., Coyle, M., Boydell, O.: Exploiting query repetition and regularity in an adaptive community-based web search engine. *User Modeling and User-Adapted Interaction* **14**(5) (2005) 383–423
38. Soller, A.: Adaptive support for distributed collaboration. In Brusilovsky, P., Kobsa, A., Nejdl, W., eds.: *The Adaptive Web: Methods and Strategies of Web Personalization*. Lecture Notes in Computer Science, Vol. 4321, Berlin Heidelberg New York, Springer-Verlag (2007) this volume.
39. Yu, Z., Zhou, X., Hao, Y., Gu, J.: TV program recommendation for multiple viewers based on user profile merging. *User Modeling and User-Adapted Interaction* **16**(1) (2006) 63–82
40. Yu, Z., Zhou, X., Zhang, D.: An adaptive in-vehicle multimedia recommender for group users. In: *Proceedings of the IEEE 62nd Semiannual Vehicular Technology Conference, Dallas (2005)*