

More Than the Sum of Its Members: Challenges for Group Recommender Systems

Anthony Jameson
DFKI, German Research Center for Artificial Intelligence
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
jameson@dfki.de

ABSTRACT

Systems that recommend items to a group of two or more users raise a number of challenging issues that are so far only partly understood. This paper identifies four of these issues and points out that they have been dealt with to only a limited extent in the group recommender systems that have been developed so far. The issues are especially important in settings where group members specify their preferences explicitly and where they are not able to engage in face-to-face interaction. We illustrate some of the solutions discussed with reference to the TRAVEL DECISION FORUM prototype. The issues concern (a) the design of suitable preference elicitation and aggregation methods, in particular nonmanipulable aggregation mechanisms; and (b) ways of making members aware of each other's preferences and motivational orientations, such as the use of animated representatives of group members.

Categories and Subject Descriptors

H.5.1 [**Information interfaces and presentation**]: Multimedia information systems—*Animations*; H.5.3 [**Information interfaces and presentation**]: Group and organisation interfaces—*Asynchronous interaction, computer-supported cooperative work, web-based interaction*

1 INTRODUCTION

Many of the items that can be recommended by recommender systems are often (or mostly) used by groups rather than by individuals, for example restaurants, vacations, movies, and TV shows. Accordingly, during the past few years a number of group recommender systems have been designed.

The most obvious additional functionality that a group recommender system has to offer, relative to a recommender for individuals, is some sort of method for aggregating (e.g., averaging) the models and/or predictions of individual users to arrive at recommendations

for the group. But as research so far has shown, the transition from an individual user to a group of users requires much more than an aggregation function.

The goal of the present paper is to formulate several of the new issues raised by group recommenders in a general way and to consider possible ways of dealing with each one, comparing these ideas with those incorporated in existing group recommender systems. Concrete illustrations will refer to a new prototype group recommender that embodies novel approaches to these general issues: the TRAVEL DECISION FORUM, which is described in an accompanying AVI 2004 system paper ([7]).

1.1 Previous Group Recommenders

It will be helpful to refer for concreteness to four of the best-known group recommender systems that have been presented so far: MUSICFX ([12]) selects music channels for the music to be played in a fitness center. On the basis of the preferences that have been previously specified by the members who are currently working out, the system chooses one of 91 possible music channels, including some randomness in the choice procedure in order to ensure variety.

LET'S BROWSE ([10]) recommends web pages to a group of two or more persons who are browsing the web together.

POLYLENS ([13]) is a generalization of the MOVIELENS system (<http://www.movielens.umn.edu/>) that recommends movies to groups of users. The system has recently been modified to yield BUDDY SEARCH, which makes it easier to form ephemeral groups.

INTRIGUE ([1]) recommends tourist attractions for heterogeneous groups of tourists that include relatively homogeneous subgroups (e.g., "children").

The recently developed TRAVEL DECISION FORUM prototype helps a group of users to agree on the desired attributes of a vacation that they are planning to take together. Special attention is given to support for users who are not collocated and who can therefore not engage in face-to-face discussions.

Since our focus is on general issues rather than specific systems, for each of these systems we will mention only the aspects that are relevant to the issues under discussion. Readers who wish to get a coherent overview of the TRAVEL DECISION FORUM are referred to the system paper ([7]) that is available elsewhere in these proceedings.

1.2 Overview of Novel Issues

Table 1 gives an overview of the issues addressed in this paper that are novel in the context of group recommendation systems (though some of them have been dealt with in other fields).

Table 1. Overview of the issues to be addressed in this paper, organized in terms of the four main phases of a group recommendation process.

Phase of the recommendation process	Difference from recommendation to individuals	Novel issue
1. Members specify their preferences.	It may be desirable for members 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 aggregating preferences must be applied.	How can the aggregation procedure effectively discourage manipulative preference specification?
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. Members decide 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 support the process of arriving at a final decision when members cannot engage in face-to-face discussion?

2 COLLABORATIVE PREFERENCE SPECIFICATION

Many recommender systems require users to specify their relevant preferences in one way or another. For example, in MUSICFX, each user rates each of 91 music genres on a five-point scale ranging from “I hate this music” to “I love this music”. The TRAVEL DECISION FORUM elicits information from each member about his or her preferences concerning several evaluation dimensions of vacations (e.g., *sports facilities*).

In an individual recommender, there is no person besides the user who has an immediate interest in seeing these 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. 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 that the current member lacks. For example, if a MUSICFX user notices that his friend has expressed a strong preference for Hawaiian music, he may decide to give it a try himself.

To exploit these potential benefits, we developed for the TRAVEL DECISION FORUM a simple extension of a typical rating-scale dialog box that allows the current member optionally to view (and perhaps copy) the preferences already specified by other members (see Figure 1).¹

A second 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.² Arguments can have various forms and functions in group decision contexts (cf., e.g., [9]). In a group recommendation context, two typical functions are

¹An earlier version of this interface was iteratively designed and tested by Claudia Plua (cf. [14]).

²These arguments can be entered and viewed in pop-up windows that are not visible in Figure 1.

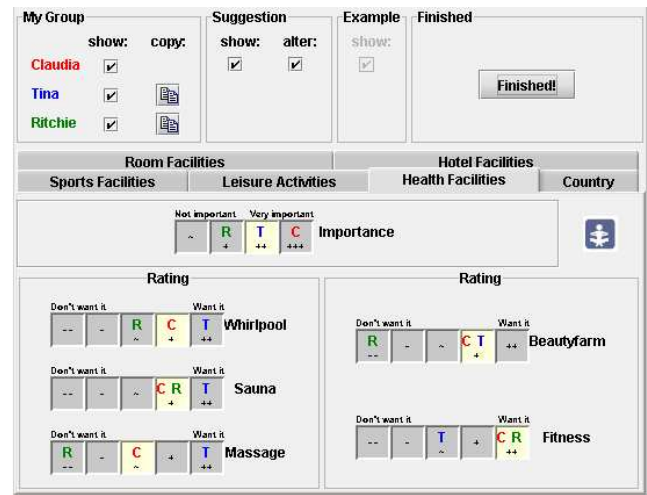


Figure 1. Dialog box for the collaborative specification of preferences.

(The currently active group member is Claudia. The preferences of each member are represented by a uniquely colored letter. 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 the mediator has suggested a compromise proposal, as is explained in Section 3.1.)

(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:

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 brief evaluation study of the collaborative preference specification interface, a typical screen shot (without verbal arguments) was presented to a group of 22 subjects, all technically oriented students. They were asked to fill in their preferences regarding hotel accommodations (a) independently and (b) while seeing the preferences of another hypothetical group member. Of the 22 subjects, 14 stated that they would prefer to see other members' preferences while specifying their own, while 5 found both options equally attractive. The subjects tended to assimilate their preferences to those of the hypothetical other member; their comments indicated that they wanted to minimize unnecessary differences in preference specifications so as to facilitate the reaching of agreement.

On the other hand, 3 of the subjects in this study preferred not to see anyone else's preferences, saying that they preferred not to be "biased" or "distracted". This overall pattern has been confirmed more informally during numerous system demos. It therefore appears advisable to make the display of other members' preferences optional, as is done in the TRAVEL DECISION FORUM.

One potential drawback of the revelation of other members' preferences concerns the possibility of *manipulative* preference specification. For example, suppose that in Figure 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 Ritchie and Tina.³

In the survey study summarized above, very few subjects showed an inclination to manipulate their preference specifications in the way just sketched; nor did they suspect that friends with whom they might want to go on vacation would be inclined to do so. But most of them did state that they would be concerned about manipulability in a setting in which there was less acquaintance and trust among the group members. It therefore appears that methods for preventing or discouraging manipulation are worth considering in connection with at least some group recommenders. This problem will be taken up in the next section.

³If she specifies her preference sincerely, the overall set of preferences in the group concerning saunas will be more positive than negative; as a result, a solution may ultimately be chosen that includes a sauna, perhaps at the expense of an attribute that Claudia does want to have.

3 NONMANIPULABLE PREFERENCE AGGREGATION

3.1 Approaches to Preference Aggregation

Even if a group recommender does not elicit members' preferences explicitly, it must have some information concerning the various users' preferences, as well as some way of aggregating this information to arrive at recommendations for the group. The topic of preference aggregation is a multifaceted and complex one that has been addressed in various scientific fields (see, e.g., [2], for a seminal contribution). Not surprisingly, then, the designers of each of the four previous systems introduced above have offered interesting discussions of alternative aggregation schemes (see, e.g., [3] and [11] for further relevant contributions).

The TRAVEL DECISION FORUM differs from the previous systems in that the main goal of the interaction is for the group members to agree on a *joint preference model*: a single way of filling out the preference specification form that can be used as a representation of the preferences of the group as a whole. That is, what the system recommends is not specific vacation solutions (e.g., concerning particular hotels in particular countries) but rather particular joint preference models, one for each value dimension. As is illustrated in Figure 2 of the overview of the TRAVEL DECISION FORUM in these proceedings ([7]), when the system has computed a recommended way of filling out the form for a given value dimension, this solution is presented to the current group member by an animated character called the *mediator*. (The roles of the other two characters visible in the figure will be explained below.) In the figure, the solution is shown both on the screen behind the mediator and in the preference specification form, through highlighting. (The full preference specification form can be seen in Figure 1.) The recommendation problem is viewed as having been solved once such a model has been agreed upon by all group members for each value dimension (cf. Section 5). This way of viewing the recommendation task is most obviously applicable when the set of specific candidate solutions (e.g., next winter's vacation catalog) is not yet available. But even when the set of possible specific solutions is already known, attempting to arrive at a common set of evaluation criteria can be an effective approach, in that it focuses attention on important differences in preferences and offers opportunities to resolve them.

3.2 Aggregation Mechanisms That Discourage or Prevent Manipulation

Here, we will concentrate on a single issue concerning preference aggregation that has received very little attention in connection with group recommenders: the problem of manipulability that was introduced in the previous section.

An early version of MUSICFX used an aggregation formula that was easily manipulable: Any music genre that was "hated" by any member currently in the fitness center was removed from the list of possible genres to play. Some users were observed to force an immediate change of music channel by adapting their specifications to indicate that they "hated" the genre currently being played.

As we saw in the previous section, manipulation can be very easy in the TRAVEL DECISION FORUM if a mechanism like averaging is used for the generation of recommended solutions. Therefore, the mediator is equipped with two types of aggregation mechanism that are *nonmanipulable* in the following sense: No group member can expect to bring about a recommendation that is more favorable

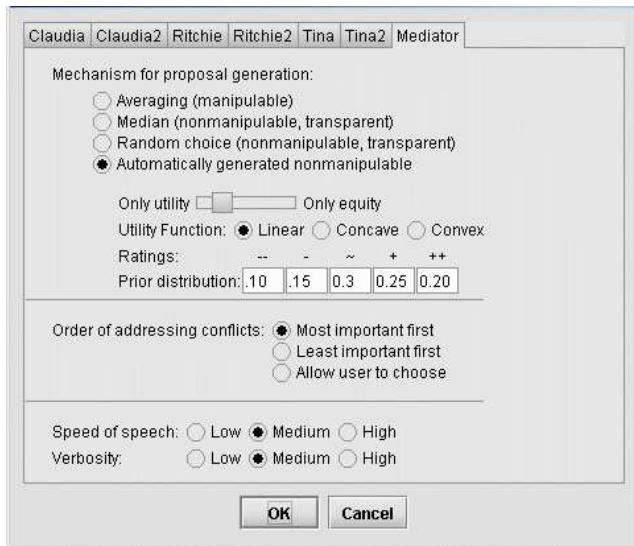


Figure 2. Dialog box for specifying the mechanism to be applied by the mediator in generating proposals.

(In the upper half, a mechanism for aggregating the preferences relative to each attribute is specified, as is explained in the text.)

to her by stating her preferences insincerely.⁴

As Figure 2 illustrates, the mediator can be instructed to use either of two types of nonmanipulable mechanisms:

1. *Hand-crafted, transparent mechanisms.* These are general aggregation mechanisms that are so simple that just about anyone can understand (a) how the mechanism works and (b) why it is nonmanipulable. For example, with the *median* mechanism, if there are three group members, the second highest preference is chosen for each attribute. Hence no member can distort the recommendation in her own favor by specifying an artificially high or low preference. This is the mechanism that was used to generate the proposal shown in Figure 1.
2. *Automatically designed mechanisms.* The hand-crafted, transparent mechanisms sometimes yield proposals that are sub-optimal in terms of overall acceptability and/or equity. For example, in Figure 1, the proposed rating for the attribute *Beauty Farm* fails to take into account the fact that Ritchie expressed a strong negative preference regarding that attribute. In an effort to create more optimal nonmanipulable mechanisms, Conitzer and Sandholm have introduced the approach of *automated mechanism design* ([4]; [5]): For a given setting, their method generates a mechanism which is (a) nonmanipulable and (b) optimal in terms of the expected value of a given *objective function* such as overall group utility (or a combination of utility and equity). Unlike the hand-crafted mechanisms, an automatically designed mechanism can also take into account specific features of a given setting such as the prior probabilities of the various possible preferences.

We investigated the applicability of automated mechanism design for group recommender systems, using the TRAVEL DECISION FORUM as an example, by (a) looking at the specific mechanisms that it generates given particular parameter settings; and (b) eliciting informal feedback about the comprehensibility and acceptability of

⁴For more complete and formal definitions, which distinguish stronger and weaker forms of nonmanipulability, see, for example, [4].

the proposals generated. The results of this study, which are reported in detail in [8], can be summarized as follows:

1. If no special measures are taken to ensure acceptability to human users, the mechanisms generated are often problematic: In particular, a mechanism is likely to be asymmetric and non-deterministic, sometimes yielding results that strike the typical user as being bizarre (e.g., a proposed joint preference that is lower or higher than any preference expressed by one of the group members). Fortunately, it is fairly straightforward within the automated mechanism design framework to impose some of the relevant constraints on the nature of the mechanisms.
2. Even if the mechanisms are designed with acceptability in mind, they tend to be more difficult than the hand-crafted mechanisms for people to understand, remember, and apply.
3. Consequently, the designer should consider carefully whether it is better to use an optimal automatically designed mechanism or a hand-crafted, transparent one—or even a familiar manipulable mechanism such as averaging. Straightforward quantitative simulations can show how much expected optimality needs to be sacrificed for the sake of familiarity and transparency.

4 VISUALIZING THE ACCEPTABILITY OF A RECOMMENDATION TO THE VARIOUS GROUP MEMBERS

Many recommender systems 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., [6]). 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.

LET'S BROWSE explains each of its web page recommendations by listing the aspects of the page that it predicts will be of interest to all group members; the system does not, however, reveal directly whether the page is predicted to be more interesting for one member than for another.

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. 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 which encourages the development of additional methods of this general sort.

INTRIGUE offers two main types of explanation: The first type presents a separate ordered list of recommended tourist attractions for each homogeneous subgroup of the entire (heterogeneous) tourist group. The second type presents a single list of recommendations in which each recommended attraction is associated with remarks about why it is suitable for particular subgroups, for example: "For yourself, [Palazzo Carignano] has high historical value. For the kids, the visit is quite short". It appears that negative aspects of an

attraction for a particular group are not explicitly mentioned. The TRAVEL DECISION FORUM introduces two novel, complementary methods that aim to provide much a more detailed picture of the consequences of a given proposal for each group member. The first method automatically follows from the use of the preference specification form for the presentation of proposals (see Figure 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., “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. Despite this wealth of information, we expected that the preference forms might be a poor substitute for the sort of feedback that group members get from each other when discussing proposals face-to-face. In such a context, a member who is disappointed with a proposal may complain about specific aspects of it in an emotional manner, formulating (or repeating) arguments. This type of reaction can heighten the group members’ awareness of each other’s points of view and overcome the natural tendency to focus on one’s own evaluations.

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. For example, in Figure 2 of the overview of the TRAVEL DECISION FORUM in these proceedings ([7]), the two animated characters on the right are the representatives of the real group members Ritchie and Tina, respectively; the current user, Claudia, is represented only by a minimal character through which she can influence the course of the interaction. Whenever the mediator has recommended a particular joint preference model for a given value dimension, he asks the representatives of the absent group members to comment on it in turn. Some typical performances of the representatives are shown in Figure 3.⁵ Each performance is generated fairly straightforwardly on the basis of the information contained in the preference specification forms; some aspects of the performances are controlled via various parameters, such as the degree of verbosity, only some of which can be explained here.

The animated characters complement the preference specification forms in the following ways:

1. *Selectivity.* In the less verbose modes, the characters verbalize only the most important aspects of the corresponding real group member’s likely response to a proposal. They therefore spare the current group member the need to extract the most important information from the preference specification forms.
2. *Vividness and familiarity.* People have much more experience in interpreting verbal expressions of (dis)satisfaction accompanied by gestures and facial expressions than they have in interpreting patterns in graphical preference specification forms.

⁵The comments of the representatives about accepting or rejecting proposals will be explained in Section 5.



Figure 3. Snapshots of some reactions of the representatives of Tina and Ritchie to the proposal shown in Figure 1.

(Tina’s representative evaluates the proposal in terms of its deviations from her own preferences: First she mentions the exact correspondences, and then she complains about the deviations. In a more verbose mode, she would also mention at the relevant points arguments specified by Tina, such as “I need a massage at the end of a strenuous day”. Finally, Tina’s representative rejects the proposal. Ritchie’s representative, by contrast, comments only on whether the proposal is in some respects better for Tina than for Ritchie.)

3. *Taking into account the represented member’s own motivational orientation.* A group member’s evaluation of a proposal does not always depend only on how well it matches his own preferences. For example, a member may be genuinely interested in maximizing the satisfaction of the group as a whole, even at the expense of his own interests. Ritchie’s representative in Figure 3 shows yet another motivational orientation: Ritchie is mainly interested in ensuring that Tina does not get a better deal than he does. How and why these motivational orientations are specified will be discussed in the next section; suffice it to say for now that there may be little point in thinking about what a proposal means for a given group member if you do not know his motivational orienta-

tion.

Since the preference forms and the representatives' performances are visible at the same time, and since the representatives use synthesized speech, the user can shift her attention between the two complementary types of presentation flexibly and to a certain extent attend to both of them at once. Moreover, the parameters governing the representatives' performances can be changed at almost any time, and the user can skip to the end of a representative's performance if she feels it is taking too long.

So far, our evaluation of people's reactions to this use of animated characters has been informal, consisting of discussions with visitors at demonstrations and a pilot study of 5 potential users who wrote down their evaluations on a questionnaire. Still, it is already clear that the performances of the animated representatives must be made an optional feature which members can use selectively or turn off completely. The feedback has confirmed that many people's first reaction to any animated characters is to turn them off on the grounds that they are distracting and/or time-consuming. Skeptical attitudes that are evidently based on previous experience with animated characters appear to be more important than any specific deficiencies of our realization of the characters (although improvements are doubtless possible).

On the positive side, a minority of those who have given feedback have expressed some appreciation for the potential advantages listed above, as well as a willingness to attend to the characters at least for a time, while they are getting used to the novel aspects of the system.

5 HELPING GROUP MEMBERS TO ARRIVE AT A FINAL DECISION

No matter how appropriate and compelling a system's recommendations and explanations are, 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.

Previous group recommender systems have tended to avoid the issue of final decision making in various ways:

1. The system simply translates the most highly rated solution into action without requesting the consent of any users.

This method is applied by MUSICFX, which switches music channels autonomously on the basis of the preferences of the group members who are present.

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. Similarly, INTRIGUE and POLYLENS appear to presuppose that one group leader will make the decision; in the field trial of POLYLENS, it was actually found that 80% of all group members requested group recommendations.

3. It is assumed that group members will arrive at the final decision through straightforward face-to-face discussion.

This assumption may be reasonable if all group members are in the same place when the system makes and explains its recommendations. Even in this relatively favorable case, effective and accurate representation of the consequences of par-

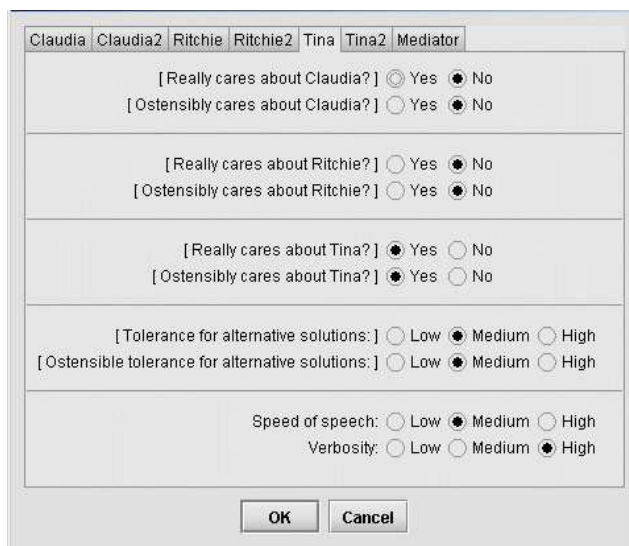


Figure 4. A dialog box for specifying the evaluation criteria of a representative concerning the absolute utility of proposals.

(Tina's representative has been instructed to evaluate a proposal solely in terms of how well it corresponds with Tina's own preferences. The representative could also be instructed to verbalize the evaluations as if Tina were more concerned about the other group members (via the options "Ostensibly cares about..."); this possibility is not discussed in this paper.)

ticular solutions for individual users can help to streamline the decision making process.

Most of the novel aspects of the TRAVEL DECISION FORUM were introduced as support for situations where group members are in different places and cannot even communicate through synchronous media such as telephones or chat rooms. Since extensive debate and negotiation via asynchronous channels such as email can be cumbersome and time-consuming, it becomes especially important to enhance mutual awareness through methods such as those described in the previous sections.

One further aspect of the TRAVEL DECISION FORUM's functionality is specifically intended to minimize the need for direct communication about recommendations: The animated representatives of absent group members (Section 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, Figure 4 shows the dialog box in which the real group member Tina has specified that her representative may accept proposals that deviate by a "medium" amount from the proposal that would be ideal from Tina's point of view. This form also shows how Tina has explicitly stated that only her own preferences are to be taken into account in the evaluation of a proposal. Figure 5 shows that Ritchie has instructed his representative somewhat differently: A proposal is to be considered undesirable to the extent to which it is more favorable for Tina than for Ritchie.

At any given moment, as was mentioned above, only one real group member is interacting with the system. If a proposal made by the mediator is accepted by the two representatives and by the current user, the mediator records that proposal as the accepted solution for the value dimension in question and proceeds to make a proposal for another value dimension. If either the current user or one of the

Claudia Claudia2 Ritchie Ritchie2 Tina Tina2 Mediator
 [Should Claudia really be better off than Tina?]
 Better Worse Equal Don't care
 [Should Claudia ostensibly be better off than Ritchie?]
 Better Worse Equal Don't care

 [Should Claudia really be better off than Ritchie?]
 Better Worse Equal Don't Care
 [Should Claudia ostensibly be better off than Ritchie?]
 Better Worse Equal Don't Care

 [Should Tina really be better off than Ritchie?]
 Better Worse Equal Don't Care
 [Should Tina ostensibly be better off than Ritchie?]
 Better Worse Equal Don't Care
 OK Cancel

Figure 5. A dialog box for specifying the motivation of a representative concerning the relative utilities of proposals for the different group members.

(Ritchie's representative has been instructed to evaluate a proposal negatively to the extent to which it is more favorable for Tina than for Ritchie.)

representatives does not accept the mediator's proposal, the current user can still try to achieve agreement on the value dimension in question by (a) changing her own preferences, perhaps in view of the positions and arguments of the other members; or (b) making a proposal of her own that may prove more acceptable to the representatives than the proposal made by the mediator.

Any value dimension upon which agreement can be reached in this way is one that will not need to be discussed (through inconvenient communication channels) by the real group members.

How well this particular method of minimizing the need for direct communication will work in any practical setting is still an open question. For now, this approach is best seen as way of calling attention to the general problem addressed in this section, which seems likely to require considerable further research as group recommenders are used with increasing frequency by noncollocated groups.

6 CONCLUSIONS

One contribution of this paper has been to formulate four challenges for the design of group recommender systems, demonstrating that they have not yet been dealt with fully in connection with the group recommenders that have been presented so far (in part, simply because they did not arise in the context of these systems).

The TRAVEL DECISION FORUM system has served here as a way of concretely illustrating some possible methods for dealing with the issues raised. The simpler methods, involving the collaborative preference specification forms, have been consistently positively evaluated both in our survey study (Section 2) and in numerous interviews, observations, and discussions. The other aspects, involving nonmanipulable aggregation mechanisms and the use of representatives of absent group members, appear to have a more limited range of applicability. The job of finding alternative approaches to the same issues is a challenge for research in this area.

7 ACKNOWLEDGMENTS

This research was supported by the German Ministry of Education and Research (BMB+F) under grant 01 IW 001 (project MIAU). Stephan Baldes, Thomas Kleinbauer, and Claudia Plua made significant contributions to the design and implementation of the TRAVEL DECISION FORUM.

8 REFERENCES

- [1] Liliana Ardissono, Anna Goy, Giovanna Petrone, Marino Segnan, and Pietro Torasso. INTRIGUE: Personalized recommendation of tourist attractions for desktop and handset devices. *Applied Artificial Intelligence*, 17(8-9):687–714, 2003.
- [2] Kenneth J. Arrow. *Social Choice and Individual Values*. Wiley, New York, 2nd edition, 1963.
- [3] Mathias Bauer and Dietmar Dengler. Group decision making through mediated discussions. In Francesco Ricci and Barry Smyth, editors, *Proceedings of the AH 2002 Workshop on Recommendation and Personalization in eCommerce*, pages 10–19. Universidad de Málaga, Departamento de Lenguajes y Ciencias de la Computación, Málaga, Spain, 2002.
- [4] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In Adnan Darwiche and Nir Friedman, editors, *Uncertainty in Artificial Intelligence: Proceedings of the Eighteenth Conference*, pages 103–110. Morgan Kaufmann, San Francisco, 2002.
- [5] Vincent Conitzer and Tuomas Sandholm. Applications of automated mechanism design. Unpublished manuscript, available from <http://www-2.cs.cmu.edu/~conitzer/>, 2003.
- [6] Jonathan L. Herlocker, Joseph A. Konstan, and John Riedl. Explaining collaborative filtering recommendations. In *Proceedings of the 2000 Conference on Computer-Supported Cooperative Work*, pages 241–250, Philadelphia, PA, 2000.
- [7] Anthony Jameson, Stephan Baldes, and Thomas Kleinbauer. 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.
- [8] Anthony Jameson, Christopher Hackl, and Thomas Kleinbauer. 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.
- [9] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.
- [10] Henry Lieberman, Neil W. Van Dyke, and Adrian S. Vivacqua. Let's Browse: A collaborative Web browsing agent. In Mark T. Maybury, editor, *IUI99: International Conference on Intelligent User Interfaces*, pages 65–68. ACM, New York, 1999.
- [11] Judith Masthoff. Group modeling: Selecting a sequence of television items to suit a group of viewers. *User Modeling and User-Adapted Interaction*, 2004. In press.
- [12] Joseph F. McCarthy and Theodore D. Anagnost. MusicFX: An arbiter of group preferences for computer supported collaborative workouts. In *Proceedings of the 1998 Conference on Computer-Supported Cooperative Work*, pages 363–372, 1998.
- [13] Mark O'Connor, Dan Cosley, Joseph Konstan, and John Riedl. PolyLens: A recommender system for groups of users. In *Proceedings of the European Conference on Computer-Supported Cooperative Work*, 2001.
- [14] Claudia Plua and Anthony Jameson. Collaborative preference elicitation in a group travel recommender system. In Francesco Ricci and Barry Smyth, editors, *Proceedings of the AH 2002 Workshop on Recommendation and Personalization in eCommerce*, pages 148–154. Universidad de Málaga, Departamento de Lenguajes y Ciencias de la Computación, Málaga, Spain, 2002.