Enhancing Mutual Awareness in Group Recommender Systems

Anthony Jameson, Stephan Baldes, and Thomas Kleinbauer *
DFKI, German Research Center for Artificial Intelligence
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
<first name>.<last name>@dfki.de

Abstract

An increasingly important type of recommender system comprises those that generate recommendations for groups rather than for individuals. The decision of a group member whether or not to accept a given recommendation can depend not only on her own evaluation of the content of the recommendation but also on her beliefs about the evaluations of the other group members and about their motivation (e.g., egocentric vs. cooperative). Yet this type of mutual awareness may be hard to achieve when group recommendations are delivered by a web-based system to group members who cannot communicate face-to-face. After introducing these issues on a general level, we discuss them more concretely by discussing a prototype group recommender system that uses several novel methods to enhance mutual awareness among group members, ranging from a group-oriented interface technique for specifying preferences to animated characters that serve as representatives of group members who are not currently available for communication.

1 Introduction

One important form of web personalization is the recommendation of products, documents, or services. The present paper does not aim to contribute a new and better solution to one of the well-known problems in this area. Instead, its main contribution is to call attention to a new opportunity (and requirement) for personalization that arises when recommendations are made not to an individual but to a group.

Even with recommendations to individuals, it is being recognized increasingly that the recommendation process does not stop once a suitable recommendation has been made. The user will have to decide which recommendation(s) (if any) to follow. To make this decision, she may need to understand the reasons for the recommendation; hence the current interest in ways of explaining recommendations (see, e.g., [Herlocker *et al.*, 2000]). She may also need information on *how* to follow the recommendation (e.g., where and how to order a particular product). In some cases, she may need information that helps her to persuade some other person (e.g., a superior who must authorize a purchase) that following the recommendation is a good idea. A recommender system is typically in a position to support at least some of these *post-recommendation processes*; and it would be short-sighted not to take these processes into account in the design of the system.

1.1 Synchronous Post-Recommendation Communication

Post-recommendation processes are even more important and complex when recommendations are being made to a group. Unless the decision as to which recommendation(s) to follow has been delegated to one group member (or to some other person), arriving at this decision will require some sort of communication and negotiation among the group members.

Consider, for example, the system MUSICFX ([McCarthy and Anagnost, 1998]), which selects a genre of music for a fitness center on the basis of the stored music preferences of the members who are currently working out. Since the selected music is not recommended but rather simply played automatically, there is no opportunity for the people present to debate about whether to accept the selection (though members have been observed trying to manipulate the selection by revising their preference specifications in artificial ways). If, instead, the training center personnel first asked the members present to approve of the system's choice of a genre (e.g., "classic rock"), one can imagine the sort of discussion that might arise. When a system's recommendations involve higher stakes (e.g., products or services that will be used by all group members and that require a significant investment), leaving the group members out of the loop is unlikely to be an acceptable option.

The easiest case is perhaps the one where the group members can communicate face-to-face (for example, if they are jointly using a single device that is delivering the recommendations): Even if the system does not provide any particular support for the discussion of the recommendations, the group members may be able to negotiate fairly effectively because of the high communication bandwidth. Each member can comment on a recommendation, stating how acceptable she finds it and perhaps why she prefers an alter-

*This research was supported by the German Ministry of Education and Research (BMB+F) under grant 01 IW 001 (project MIAU).

native solution. These comments may cause other members to reconsider and revise their own specified preferences, especially if the preferences were not initially specified very carefully or completely. The members may refer to group-related considerations such as equity (e.g., "This solution is much better for the two of you than it is for me") and overall group utility. They may become aware of differences in each others' overall orientation; for example, it may become clear that Member A is thinking only of his own interests whereas Member B is thinking about how to ensure that all members of the group end up being reasonably happy with the chosen solution. In fact, a group member's responses to the recommendations made by the system may be influenced more strongly by such group-related considerations than by the member's own evaluation of the recommendations (see, e.g., [Carnevale and Pruitt, 1992] for a discussion of the importance of such motivational orientations and the associated strategies in negotiations). In short, communication and mutual awareness play important roles when members of a group need to arrive at a joint decision, even if it involves a simple selection from a set of recommended solutions.

1.2 Asynchronous Post-Recommendation Communication

Consider now a more difficult scenario, which may be more typical of web-based group recommender systems: The group members are not considering the recommendations in the same place or even at the same time, and they are not in a position to use synchronous communication media like the telephone, audio- or videoconferencing, or text chat. How are the group members supposed to arrive at a satisfactory decision? More specifically, how can we recreate at least some of the broad-bandwidth communication and mutual awareness that figure so importantly in the face-to-face situation? Traditional asynchronous communication media like email can convey some of the relevant information (e.g., each group member can write messages commenting on the recommendations made). But in addition to being laborious, this type of communication is wasteful in that much of the information conveyed is information that is already possessed to some extent by the recommender system: In order to be able to generate recommendations, the system must know something about the members' preferences, and it must be able to predict members' evaluations of particular solutions.[1] A system should therefore be able to create some degree of mutual awareness while limiting the need for group members to invest time and effort in post-recommendation communication.

1.3 Relevant Previous Work

In the field of computer-supported cooperative work, many techniques have been developed for enhancing awareness of other users (see, e.g., [Liechti, 2000] for an overview of such technique for web-based systems). Although most of these techniques concern forms of awareness that are not especially applicable to group recommender systems, some ideas can be adopted. For example, the collaborative editor PREP (see, e.g., [Neuwirth *et al.*, 2000]) uses separate columns to display the comments of the several authors of a jointly authored document; this method was one source of inspiration for the collaborative preference specification method discussed in 2.1.

One particularly suggestive group recommender system is POLYLENS ([O'Connor *et al.*, 2001]), which uses collaborative filtering to recommend movies to a group of acquaintances who want to go to a movie together. The authors discuss various possible ways of displaying recommendations; the method they chose for POLYLENS shows the predicted rating of each recommended movie for the group as a whole and for each of the group members individually. This method allows users to see, for example, how satisfied with each movie the least satisfied group member is likely to be. Although this information could be of use in a post-recommendation discussion among group members, the authors do not discuss such post-recommendation processes. (Since people who are planning to see a movie together are usually located near each other, asynchronous computer-supported communication is presumably seldom necessary.)

A characteristic feature of the prototype to be presented here is the use of animated characters that serve as representatives of group members who are not currently on-line. Animated characters have often been used for personalized recommendation, but the characters have in general represented persons doing the recommendation, and the emphasis has been on effective communication about the recommended items (see, e.g., [Cassell, 2001]). André and Rist [2001] introduced the concept of *teams* of product presenters. Their characters exhibit personality and express emotional reactions to aspects of the products under discussion. The present work makes use of much of the same technology, but the main role of the animated characters is different: to serve as representatives of absent group members and thereby to enhance awareness of their points of view.

1.4 Overview of the Travel Decision Forum Prototype

In the main part of this paper, we will discuss the issues raised above more concretely by introducing a prototype system that uses several methods to enhance mutual awareness in a group recommendation context. The goal is not to propose this particular prototype as an implementation solution but rather to suggest ideas that can be adapted in different ways, in a variety of contexts.

The prototype supports users in the following scenario: Three friends want to plan a joint vacation. They are not able to get together or to use synchronous communication media to discuss their plans. Consequently, at any moment only one group member will be interacting with the web-based system, which will make use of stored information about the other group members.

Figure 1 shows the setting in which the interaction takes place: An animated character, the *mediator*, sits in front of

[1] Exactly what information of this sort the system possesses depends on the nature of the recommendation method(s) used. The present paper considers the type of system for which users specify their preferences explicitly. The methods required, say, for a pure collaborative filtering recommender would have to be in part quite different.

Figure 1. *Snapshot of an interaction in the Travel Decision Forum.*
The mediator's proposal for the dimension Health Facilities *is shown on the screen, as well as in the preference form at the bottom left (shown fully in Figure 3). The two representatives have just rejected the proposal (cf. Figure 4 below), and the current user Claudia must now decide how to respond to it herself.*

a screen on which he can display possible solutions and proposals. On the right, we see two animated characters that represent the two absent group members (here: "Ritchie" and "Tina"). In the front, we see the back of a character that represents the currently active group member (here: "Claudia"). Unlike the other two characters, this character does not act independently but rather serves as a means of making the current member visible on the screen. In this snapshot, Claudia has the opportunity to respond to a proposal made by the mediator, having heard the responses of the representatives of Ritchie and Tina. We will now describe the interaction that leads up to scenes like the one depicted in the figure.

2 Phase 1: Specification and Refinement of Preferences

2.1 Collaborative Preference Specification

In the first of the two main phases of the interaction with the system, each member specifies her preferences concerning the vacation via a collaborative preference specification form (Figure 2). This form is in many ways similar to preference specification forms that are familiar from systems like the ACTIVE SALES ASSISTANT (http://www.activebuyerguide.com/) and the once-popular system PERSONALOGIC (which is no longer available): For each of several *value dimensions* (shown on the tabs in the form), a number of *attributes* are shown. For each attribute, the user can give a rating between $--$ ("Don't want it") and $++$ ("Want it"). These ratings are interpreted in terms of multiattribute utility theory (cf. [Jameson *et al.*, 1995]); we omit details concerning their meaning and processing, since these are not required for an understanding of the central points being made here. On the basis of the preference specifications of each member, the system can evaluate a specific vacation solution from the point of view of that member. Similarly, after aggregating the preferences of the three members according to any of various possible methods, the system can evaluate any specific solution from the point of view of the group as a whole.

The novel aspect of this preference specification form is the way in which it allows the previously specified preferences of other group members to be viewed (optionally) by the current member. This feature can be seen as a simple way of enhancing awareness of the preferences of other group members (cf. the discussion of an earlier version by [Plua and Jameson, 2002]). Despite its simplicity, this collaborative preference specification appears to yield considerable benefits. For example, in a brief study that we conducted with 22 subjects, 14 stated that they preferred being able to see the preferences of another group member while specifying their

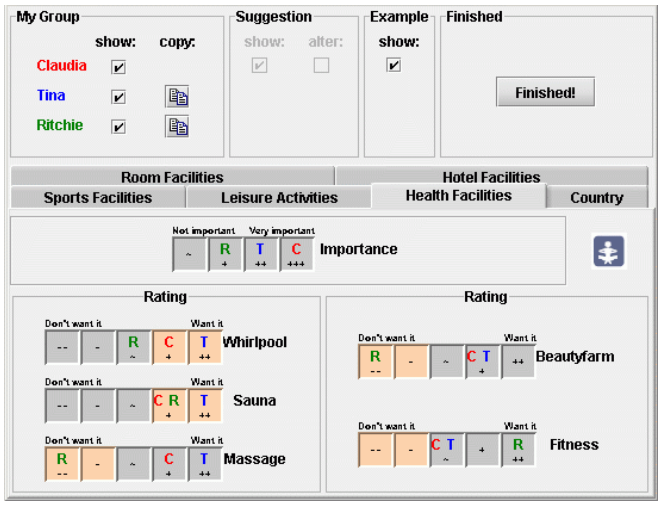

Figure 2. *Dialog box for the collaborative specification of preferences.*

The currently active group member is Claudia. The preferences of each member are represented by a uniquely colored letter; Claudia cannot change the position of the letters "T" and "R". The highlighting of two of the cells for each attribute is added only when the mediator has presented an example solution: Highlighting of the two rightmost cells in a scale indicates that the attribute in question is present in the example solution.

own, while only 3 preferred not to have them shown and 5 were indifferent. When the preferences of another member were shown, there was a tendency for the current user's specified preferences to be more similar to the other member's preferences. Subjects' comments indicated that they wanted to minimize unnecessary differences in preference specifications so as to facilitate the reaching of agreement.

2.2 Testing Preference Specifications by Requesting Example Solutions

During this first phase, at any point the current member can perform a "reality check" on her specified preferences by referring to the items in a relevant database of possible solutions. In the simplest case, the current member asks the mediator to show the solution from the database that fits best with her own preferences as they have been specified so far—that is, without regard to the preferences of the other members. It is well known from experience with recommender systems that users often want to check intermittently what kind of solutions are likely to emerge, even after they have input just a few preferences or ratings of their own. Not only does this type of intermittent testing make the interaction more interesting; it also gives the user important feedback about her preferences—for example, concerning unrealistic requirements or neglected aspects of the problem.

This type of incremental preference elicitation is supported in many recommender systems (e.g., the ACTIVE SALES ASSISTANT mentioned above, the AUTOMATED TRAVEL ASSISTANT of [Linden *et al.*, 1997], and the VEIL system of [Blythe, 2002]). Our prototype once again introduces a somewhat novel group-related element: The current member can ask to see an example solution that is based on the specified preferences of all of the group members, along with a visualization of its relationships to the preferences of the individual members. The example solutions are summarized on the screen behind the mediator and also shown in detail through marking in the preference form. For example, in Figure 2, Claudia can see that with the current example solution, as far as *Health Facilities* are concerned, her own expressed preferences are fulfilled only with respect to the attributes *Whirlpool* and *Sauna*. This result partly reflects the overall pattern of preferences of the three members, but it is also due in part to the distribution of these facilities in the hotels that fulfill the group members' other requirements. Contemplating the information shown in the form, Claudia may decide to shift her priorities for the planned vacation and adjust her preferences accordingly.

This first phase continues for the current member until that member has indicated that she is not interested in continuing to specify or adapt her preferences in this way.

3 Phase 2: Arriving at a Joint Preference Model

3.1 The Goal of This Phase

In this phase, the group members' goal is 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 mediator 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. The decision making problem is viewed as having been solved once such a model has been agreed upon for each value dimension. This way of viewing the problem is most obviously applicable when the set of specific candidate solutions (e.g., next winter's vacation catalog) is not yet available. In this case, the group members can do no more than agree on their criteria, and any specific solutions (e.g., from last winter's catalog) can serve only as examples, in the way illustrated in 2.2. But even when the set of possible specific solutions is already known, arriving through discussion at a common set of criteria can be an effective approach, in that it focuses attention on important differences in preferences and offers opportunities to resolve them.

A somewhat different—and more familiar—form of interaction would be for the mediator to recommend only specific solutions, which would then be discussed by the group members and their representatives. The methods discussed in this paper could also be applied in this type of setting, after some adaptation.

3.2 Structure of the Interaction in Phase 2

In this second phase, animated characters representing the two absent group members appear on the screen (cf. Figure 1). For each value dimension, the mediator moderates a simple form of negotiation between the current member and the two representatives, proceeding as follows:

Step 1. On the basis of the specified preferences of all members, the mediator computes a *proposal*—a particular

Figure 3. *Display of a proposal in the preference specification form.*

Here, the highlighting of one cell for each attribute represents the proposed joint rating for that attribute.

way of filling out the preference specification form with respect to one value dimension. (Methods for generating proposals will be discussed in 3.5.) The mediator then displays the proposal on the screen behind him. The screen in Figure 1 shows a proposal for the value dimension *Health Facilities*. This proposal is also displayed in the preference specification form, which the current member may inspect at any time: For each attribute, the proposed joint rating is highlighted with a white background (see Figure 3).

Step 2. The mediator gives each of the two representatives an opportunity to portray the reaction to this proposal of the corresponding real group member. For example, Figure 4 shows some of the reactions of the two representatives to the proposal shown in Figure 3 (see 3.3 for further discussion). Each performance ends with a statement of acceptance or rejection of the proposal, which depends on whether the representative's evaluation of the proposal exceeds a threshold specified by the corresponding real group member.

Step 3. The mediator asks the current member to respond to the proposal. This response can take various forms:

1. The current member accepts the proposal.
   If the other representatives have likewise accepted the proposal, the mediator declares the discussion of the current value dimension to be completed, and the proposal is stored as part of the joint preference model.[2]
2. The current member adapts her preferences.
   The performances of the two representatives may, for various reasons, have induced the current member to reconsider and change her own preferences. In this case, the mediator will present in the next step a new proposal which is based on the new set of individual preference models.

[2]If one or both of the representatives has rejected the proposal, it makes little sense for the current member to accept it; one of the next two responses listed is more appropriate.

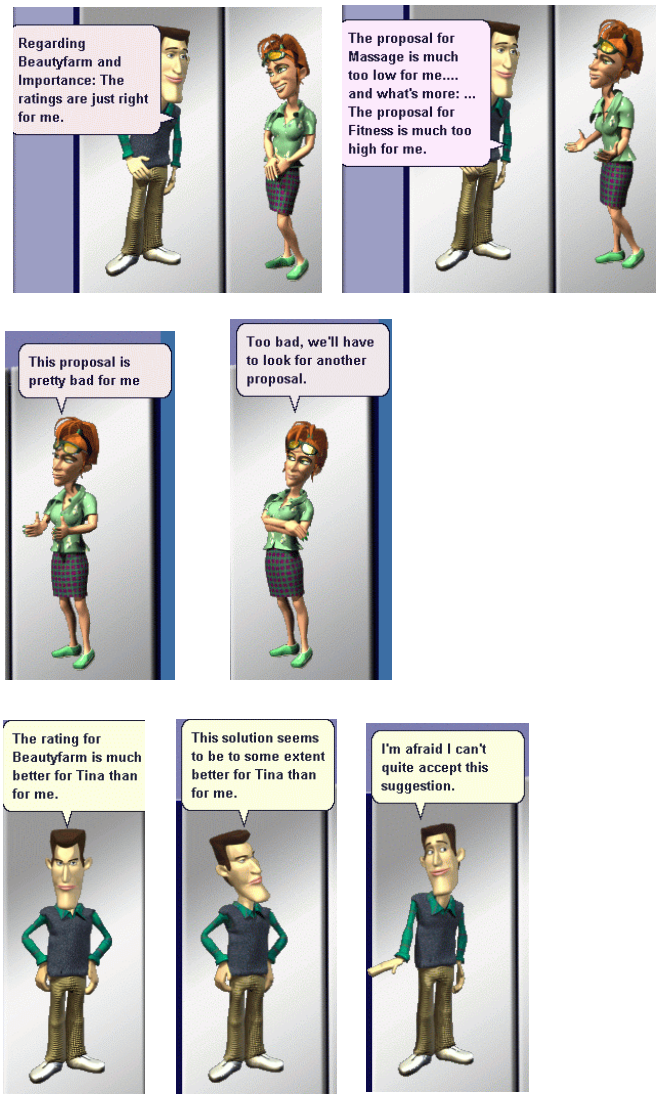

Figure 4. *Snapshots from the performances of the representatives of Tina and Ritchie to the proposal shown in Figures 1 and 3.*

Tina evaluates the proposal in terms of its deviations from her own preferences: First she mentions the exact correspondences, then she complains about the deviations, and finally she rejects the proposal. Ritchie is interested only in whether the proposal is in some respects better for Tina than for him.

3. The current member adapts the proposal.
   For example, in Figure 1, Claudia may notice that the mediator's proposal happens to correspond with her own preferences for all attributes. She can try to think of possible changes to the proposal that will make it more acceptable to Tina and Ritchie's representatives—not an easy task in this particular case.
4. The current member rejects the proposal.
   This response makes sense only when the current member sees no likelihood that further discussion of this value dimension might lead to a mutually acceptable solution. In this case, the mediator puts the current di-

mension aside for the time being. He will bring it up again in the next interaction with one of the other group members; perhaps this member will see some way of achieving consensus through a response of type 2 or 3.

Step 1 again. The mediator goes back to Step 1, presenting a new proposal for discussion. Depending on what has just happened in Step 3, this proposal will be either a counterproposal offered by the current member or a proposal generated by the mediator concerning either the current value dimension or the next dimension.

3.3 Performances of the Representatives

The performances of the animated characters (Figure 4) are designed to place the current member in a position that is more or less similar to that of a person who, in face-to-face communication with another person, tries to recognize that person's motivation, preferences, and specific points of view.

Each representative responds to a proposal by (a) commenting on the aspects of the proposal that are most important for the group member that it represents, (b) summarizing that group member's presumed overall evaluation, and (c) announcing acceptance or rejection. Gestures, facial expressions, and speech output make these performances more natural and vivid, while making it less necessary for the current member to attend to the details of the formulations.

One typical consequence of these performances is to make it clear to the current member that not only she but also the other members are in general making sacrifices if they agree to a compromise proposal. The current member's own sacrifices tend to be relatively salient to her, and they may even evoke emotional responses—for example, when she agrees to do without her favorite sports activity or (as a smoker) to stay in a nonsmoking room. It is less obvious to her, without some sort of emphasis, that the other members have to make similar sacrifices: These sacrifices can in principle be recognized in the preference specification form (cf. Figure 3) as cases where the letter for the member in question lies (far) away from the highlighted cell representing the proposal. But this representation is much more abstract than the current member's own responses; and because of the large amount of information in the form, relevant information can easily be overlooked.

The performances of a representative also convey an impression of the motivational orientation of the group member in question. On the basis of utterances like those shown in Figure 4, it is fairly easy to recognize motivational orientations like egocentrism or an obsession with equity, which would be harder to communicate in other ways.

Although the current member cannot change the preferences or motivation of the representative of another group member, she can influence the speed and degree of detail of the representatives' performances. In particular, a user can gradually diminish her own reliance on the performances of the animated characters, if she so desires, relying more and more on the information presented compactly in the preference specification form, once she has become accustomed to interpreting it.

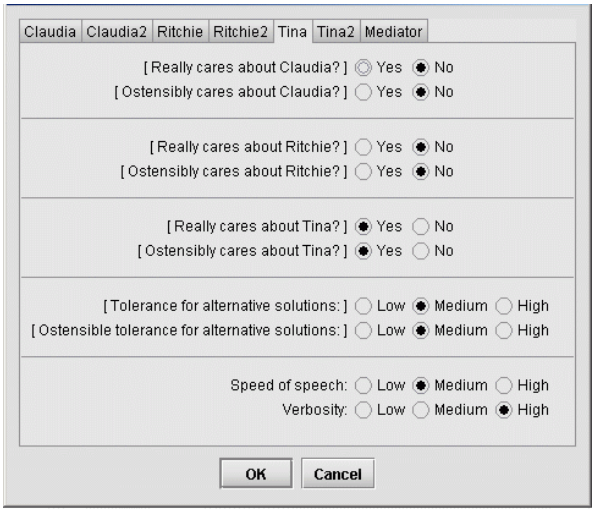

Figure 5. *A screen for specifying the evaluation criteria of a representative concerning the absolute utility of proposals.*
Tina's representative has been instructed to evaluate a proposal solely in terms of how well it corresponds with Tina's own preferences. The representative could also be instructed 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.

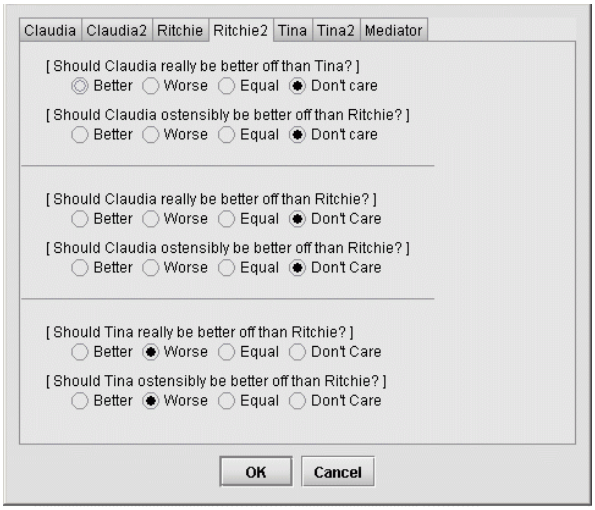

Figure 6. *A screen for specifying the motivation of a representative concerning the relative utilities of proposals for the different group members.*
Ritchie's representative has been instructed to evaluate a proposal negatively to the extent to which it is more favorable for Tina than for Ritchie.

3.4 Motivational Orientation of the Representatives

As was mentioned in 1.1, members in a decision making group can have a wide variety of motivational orientations. Accordingly, the prototype allows each real group member to specify not only her preferences but also the motivational orientation that her representative is to adopt.

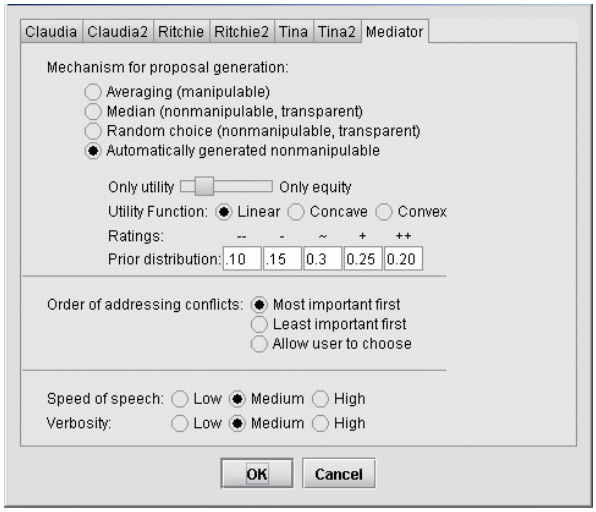

Figure 7. *Screen for specifying the strategy to be applied by the mediator in generating proposals.*
In the upper half, a mechanism for aggregating the preferences relative to each attribute is specified, as is explained in the text.

Figures 5 and 6 show the key settings for Tina and Ritchie, respectively. These two orientations, chosen here for purposes of illustration, are simple and egocentric: The difference is that Tina is concerned that a proposal should be favorable for her, whereas Ritchie is worried only about the relative desirability of a proposal for Tina vs. himself. By contrast, Tina could have specified a cooperative motivation in Figure 5 by choosing the option *Yes* with respect to all three group members. Similarly, concern about absolute utility can be combined with concern about relative utilities.

The motivational orientation of a representative determines which aspects of a given proposal a representative (a) will comment on and (b) will take into account when deciding whether to accept or reject the proposal. Specifically, if the representative is concerned about the absolute utility for a given person, it will take into account any deviations of the proposal from the preferences of that person; how large the deviation has to be in order to be worthy of comment is determined by a *verbosity* parameter that the current member can set. If the representative is concerned about relative utility, it takes into account *differences* in deviations.

3.5 Methods for Generating Proposals

When a proposal is to be generated by the mediator (instead of by the current member), part of the goal is to compute a proposal that is likely to be acceptable to the current member and to the two representatives. Although a thorough discussion of this problem is not possible here, a sketch of the approach taken in the prototype will call attention to some typical issues that arise (see also [Jameson *et al.*, 2003]).

The prototype provides several mechanisms by which the mediator can generate a proposal, the choice of which can be influenced by the mediator and/or the current member:

1. The *averaging* mechanism: For each attribute, the average of the three ratings is chosen, with the result being rounded off where necessary.
2. The *median* mechanism: For each attribute, the median of the three ratings is chosen.
3. The *random-choice* mechanism: For each attribute, one of the three ratings supplied is chosen randomly.
4. An automatically generated *nonmanipulable* mechanism.

Mechanisms of this fourth type are generated via a specific application of a general method devised by Conitzer and Sandholm [2002]: A (possibly nondeterministic) mechanism is computed on the fly that maps each possible set of three individual ratings onto a single joint rating. The generated mechanism is *nonmanipulable*: It ensures that a group member cannot expect that a proposal more favorable to her will be generated if she specifies her preferences inaccurately. This potentially desirable property of a mechanism is shared by the median and random-choice mechanisms; but with the averaging mechanism, a member can sometimes be tempted to give an inaccurate, extreme rating for a given attribute in order to ensure that the proposal generated for that attribute corresponds with her own true preference. Attempts by individual group members to gain an advantage (perhaps unconsciously) by distorting their preference specifications can be a serious issue in group recommender systems, as early experience with the MUSICFX system illustrated ([McCarthy and Anagnost, 1998, p. 370]). Measures that enhance group members' awareness of each others' preferences can actually encourage such manipulation by making it easier to determine the optimal manipulation.

As can be seen in Figure 7, within the constraint of nonmanipulability, the automatically generated mechanism is designed so as to maximize a weighted average of (a) overall utility for the three group members and (b) equity, the extent to which a proposal is equally desirable for all members. This maximization process takes into account the a priori probabilities of particular preferences. For example, the probability distribution specified in Figure 7 reflects the assumption that preferences on the left-hand side of the scale are relatively infrequent.[3]

The use of nonmanipulable mechanisms for proposal generation does not eliminate all opportunities for manipulation: Group members are not obligated to accept any of the proposals generated, and they have opportunities to change their preferences repeatedly. Moreover, automatically generated mechanisms tend to be relatively hard for users to understand and accept. Still, nonmanipulable mechanisms can contribute to efforts to deal with the problem of insincere preference specification.

3.6 Ultimate Results of the Interaction

The interaction described in this section continues until the current member has either (a) agreed (to the extent possible) with the representatives of the other members on a joint preference model for each value dimension or (b) run out of time

[3] Our tests with potential users have shown that a negative preference concerning a given attribute typically arises when one group member is worried that other group members will spend too much time on the option in question, reducing the time available for joint activities.

or interest. At this point, even if there are still some value dimensions for which no joint preference model has been agreed upon, there may have been considerable movement in the direction of an agreement concerning these dimensions: The current member may have adapted her preference specifications in a way that reduces the amount of conflict; she may have adopted a more cooperative orientation toward the evaluation of proposals; and she may have increased the leeway granted to her representative in accepting compromise proposals. When the next group member interacts with the system (seeing, of course, the animated character representative of the previous current member instead of his own representative), there will be opportunities for further convergence.

3.7 The Overall User Experience

The ways in which the methods embodied in this prototype enhance mutual awareness, especially in Phase 2, can be fully appreciated only through interaction with the prototype itself. After a brief period of getting used to the rather unusual aspects of the interaction, the current group member finds herself focusing on the higher-level question: "How can I arrive at a satisfactory agreement with these two people concerning the aspect of the decision currently under discussion?" In doing so, the current member tries to understand the representatives' reactions to previously discussed proposals, to assess their motivation and the amount of leeway that they have in accepting proposals, and to predict how they will react to new proposals. These judgments can be based on the behavior of the animated characters, on the information presented in the preference form, or (most likely) on both sources of information. The user can adjust most aspects of the behavior of the animated characters to suit her own taste.

4 Concluding Remarks

The presentation of the Travel Decision Forum has called attention to a number of issues for group-oriented recommendation that do not arise with recommendation that is oriented toward individuals. For practical settings in which communication among group members about recommendations is desirable or necessary, and in which only asynchronous communication is possible, it may be worthwhile to adapt some subset of the methods presented here. Even in settings where these methods are infeasible (e.g., because of technical limitations) or unnecessary (because adequate communication among group members can take place through more conventional channels), designers and researchers may find the prototype presented here helpful as a reminder of the importance and complexity of post-recommendation communication and decision processes.

References

[André and Rist, 2001] Elisabeth André and Thomas Rist. Presenting through performing: On the use of multiple lifelike characters in knowledge-based presentation systems. *Knowledge-Based Systems*, 14(1–2):3–13, 2001.

[Blythe, 2002] Jim Blythe. Visual exploration and incremental utility elicitation. In Rina Dechter, Michael Kearns, and Richard S. Sutton, editors, *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, pages 526–532. AAAI Press / MIT Press, Menlo Park, CA / Cambridge, MA, 2002.

[Carnevale and Pruitt, 1992] P. J. Carnevale and D. G. Pruitt. Negotiation and mediation. *Annual Review of Psychology*, 43:531–582, 1992.

[Cassell, 2001] Justine Cassell. Embodied conversational agents: Representation and intelligence in user interfaces. *AI Magazine*, 22, 2001.

[Conitzer and Sandholm, 2002] 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.

[Herlocker *et al.*, 2000] 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.

[Jameson *et al.*, 1995] Anthony Jameson, Ralph Schäfer, Joep Simons, and Thomas Weis. Adaptive provision of evaluation-oriented information: Tasks and techniques. In Chris S. Mellish, editor, *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1886–1893. Morgan Kaufmann, San Mateo, CA, 1995.

[Jameson *et al.*, 2003] Anthony Jameson, Christopher Hackl, and Thomas Kleinbauer. Evaluation of automatically designed mechanisms, 2003. Manuscript submitted for publication, available from the authors.

[Liechti, 2000] Olivier Liechti. Awareness and the WWW: An overview. *ACM SIGGROUP Bulletin*, 21(3):3–12, 2000.

[Linden *et al.*, 1997] Greg Linden, Steve Hanks, and Neal Lesh. Interactive assessment of user preference models: The automated travel assistant. In Anthony Jameson, Cécile Paris, and Carlo Tasso, editors, *User Modeling: Proceedings of the Sixth International Conference, UM97*, pages 67–78. Springer Wien New York, Vienna, 1997.

[McCarthy and Anagnost, 1998] 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.

[Neuwirth *et al.*, 2000] Christine M. Neuwirth, David S. Kaufer, Ravinder Chandhok, and James H. Morris. Issues in the design of computer support for co-authoring and commenting. In *Proceedings of the 2000 Conference on Computer-Supported Cooperative Work*, pages 183–195, 2000.

[O'Connor *et al.*, 2001] M. O'Connor, D. Cosley, J. Konstan, and J. Riedl. PolyLens:A recommender system for groups of users. In *Proceedings of the European Conference on Computer-Supported Cooperative Work*, 2001.

[Plua and Jameson, 2002] 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.