Two Methods for Enhancing Mutual Awareness in a Group Recommender System

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

We present a group recommender system for vacations that helps group members who are not able to communicate synchronously to specify their preferences collaboratively and to arrive at an agreement about an overall solution. The system's design includes two innovations in visual user interfaces: 1. An interface for collaborative preference specification offers various ways in which one group member can view and perhaps copy the previously specified preferences of other users. This interface has been found to further mutual understanding and agreement. The same interface is used by the system to display recommended solutions and to visualize the extent to which a solution satisfies the preferences of the various group members. 2. In a novel application of animated characters, each character serves as a representative of a group member who is not currently available for communication. By responding with speech, facial expressions, and gesture to proposed solutions, a representative conveys to the current real user some key aspects of the corresponding real group member's responses to a proposed solution. Taken together, these two aspects of the interface provide complementary and partly redundant means by which a group member can achieve awareness of the preferences and responses of other group members: an abstract, complete, graphical representation and a concrete, selective, human-like representation. By allowing users to choose flexibly which representation they will attend to under what circumstances, we aim to move beyond the usual debates about the relative merits of these two general types of representation.

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*
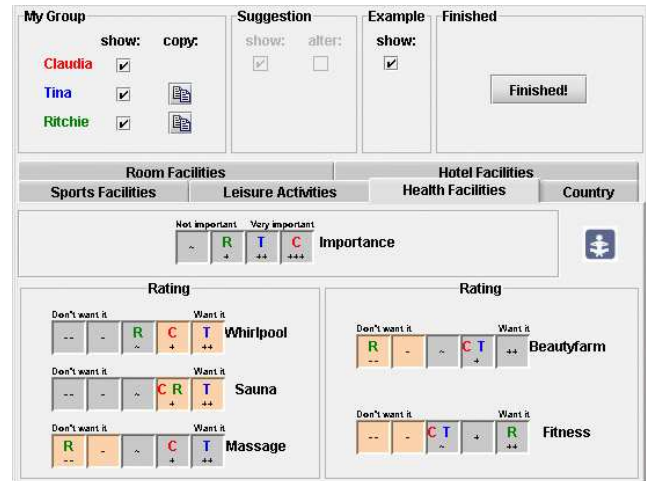

Figure 1. *Dialog box for the collaborative specification of preferences.*

(The currently active group member is Claudia. The preferences of each member are represented by a uniquely colored letter; 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: For example, highlighting of the two rightmost cells in a scale indicates that the attribute in question is present in the example solution.)

1   INTRODUCTION

This paper gives a brief overview of the TRAVEL DECISION FORUM, which will be demonstrated interactively at AVI 2004. For a fuller discussion of the theoretical considerations underlying the system, as well as a systematic comparison with other group recommender systems and other related research, the reader should consult the companion full paper ([1]), which also contains most of the figures that are referred to in this paper.

The TRAVEL DECISION FORUM helps three members of a group to agree on a single set of criteria that are to be applied in the making of a particular decision (specifically, what their planned joint vacation should be like). At any given moment, at most one member is interacting with the system. Figure 2 shows a typical snapshot of an interaction in the TRAVEL DECISION FORUM. The character on the left is the *mediator*, who directs the interaction between the current user (represented by the figure at the bottom of the screen) and the two *representatives* of the other two users (shown at the right). The elements in this screen shot will be introduced one by one below.

Figure 2. *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 1 of the long paper [1]). The two representatives have just rejected the proposal, and the current user Claudia must now decide how to respond to it herself.)*

2 INITIAL SPECIFICATION OF PREFERENCES

In an initial phase of the decision making process, each group member specifies her preferences regarding the joint vacation (at least partially) by filling in a *preference specification form* (see Figure 1). Novel aspects of this form include the possibilities of (a) optionally viewing, copying, and post-editing the preferences that have already been entered by other group members and (b) adding verbal arguments (which can be viewed by other members) to explain particular specifications. (These arguments are entered and viewed in pop-up windows that are not visible in Figure 1.)

At any time, the user can obtain feedback on the consequences of the preferences that she has specified so far by requesting *example solutions*. These are specific vacation plans that would be retrieved from the database of possible solutions, either (a) on the basis of only her own specified preferences or (b) on the basis of the preferences specified by all group members. For each example, the mediator lists the most important attributes on the screen behind

him. He also makes detailed information about the example solution available in the preference specification form by highlighting, for each attribute of the solution, the corresponding cells in the preference form, as is illustrated in Figure 1). This visualization of the solution allows the user to see at a glance which attributes of the solution fulfill (or violate) her preferences—and those of the other group members. This type of feedback gives the user a sense of the possibly orthogonal constraints imposed by (a) the availability of satisfactory solutions in the database and (b) conflicts between the user's own preferences and those of the other group members.

3 AGREEING ON A JOINT PREFERENCE MODEL

At the end of this initial phase, each member will have specified at least some preferences, and there will in general be some important differences in these preferences. The goal of the second phase is for all group members to agree on a single *joint preference model*:

a single way of filling in the preference specification form that can be used to retrieve suitable specific vacation solutions.

In this phase, when one group member interacts with the system, she sees three animated characters (see Figure 2): a *mediator*, who controls the interaction, and *representatives* of the two group members who are not currently on-line.[1] For each value dimension (e.g., *health facilities*) in turn, the mediator examines the specified preferences of the three members and recommends a single preference specification for that dimension which may represent an acceptable compromise for all members. The mediator uses one of a number of preference aggregation functions to arrive at the compromise proposals, including some that are designed to discourage the manipulative specification of preferences. The mediator's proposal is displayed on the screen behind the mediator and (through highlighting) in the preference specification form, which the current user can view at any time (see Figure 1 in the long paper, [1]). The mediator asks each of the two representatives in turn whether it can accept the compromise proposal on behalf of the corresponding absent real group member. In the simplest case, the representative makes this decision by determining how far the proposal deviates from its real member's stored preferences, checking whether the deviation falls below the threshold of maximum deviation that has been specified by that member. But the real group member may also have instructed the representative to take into account the preferences of other members (e.g., to determine the suitability of the proposal for the group as a whole).

In addition to accepting or rejecting the proposal, each representative offers a more or less detailed explanation of its decision in terms of specific aspects of the proposal that do or do not correspond with the preferences of the group member(s) whose preferences are being taken into account. Figure 3 in the long paper ([1]) shows some examples of how representatives respond to proposals. These performances of the representatives, together with the information in the preference specification form, serve to heighten the current real member's awareness of the preferences of the absent members—as well as her awareness of their motivational orientations (e.g., whether they are concerned about the group as a whole or only about their own interests). On the basis of this awareness, the current member can decide how to respond to the compromise proposal herself: If she accepts it (and if the two representatives have also accepted it), the proposal is adopted as part of the joint preference model, and attention is shifted to another value dimension. Otherwise, the current user can (a) adapt her own preference specification, perhaps under the influence of the comments made by the two representatives; or (b) suggest an alternative proposal of her own that seems more likely to be accepted by the two representatives.

When the current member has finished interacting with the system, another member can interact with it at a later time, dealing only with the value dimensions on which agreement has not yet been reached.

Direct communication among the group members, which is assumed to be costly, is required only with regard to any value dimensions on which agreement could not be reached via the procedure sketched above.

---

[1] For the implementation, the MIAU platform described in [2] was used.

4 SIGNIFICANCE FOR ADVANCED VISUAL INTERFACES

As is discussed in the companion long paper, the significance of the TRAVEL DECISION FORUM lies largely in the way in which it highlights general issues and possible design solutions that have so far received little attention in work on group recommender systems. Another contribution of this system to research on advanced visual interfaces concerns the long-standing debate about the relative merits of two contrasting interaction metaphors: powerful tools for information manipulation and visualization vs. anthropomorphic characters that imitate human-human communication (see, e.g., [3]). The TRAVEL DECISION FORUM instantiates both of these metaphors with its two methods for conveying information about the preferences and responses of other group members. We can therefore observe which style of interaction users seem to prefer—and why—without having to construct contrasting experimental conditions. So far, feedback from two small-scale studies and several public demos indicates that the simple visualizations are almost universally appreciated, whereas the animated characters evoke mixed reactions.

But the more interesting question, which can be answered only when the two contrasting approaches are combined within a single system, concerns ways in which users can use both types of information alternately and/or in parallel, deriving complementary benefits from each type. For example, a user can scan a dialog box such as the one shown Figure 1 of the long paper ([1]) to get a visual overview of conflicts and agreements with regard to all of the attributes on a particular dimension; and at the same time she can listen to an animated representative who is expostulating about a couple of points that are particularly important for one group member.

We can also observe the extent to which users' attention gradually shifts from one type of system output to the other as they gain experience in the system. In particular, a novice user may appreciate the immediate comprehensibility of the performances of the representatives; but as she becomes more skilled at recognizing relevant visual patterns, she may come to rely increasingly on the visualizations and ultimately fade the representatives out entirely.

5 ACKNOWLEDGMENTS

6 REFERENCES

[1] Anthony Jameson. 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.

[2] Thomas Rist, Elisabeth André, and Stephan Baldes. A flexible platform for building applications with life-like characters. In W. L. Johnson and Elisabeth André, editors, *IUI 2003: International Conference on Intelligent User Interfaces*, pages 158–165. ACM, New York, 2003.

[3] Ben Shneiderman and Pattie Maes. Direct manipulation vs. interface agents. *interactions*, 4(6):42–61, 1997.