

Group Recommender Systems: A Critiquing Based Approach*

Kevin McCarthy, Maria Salamó, Lorcan Coyle, Lorraine McGinty,
Barry Smyth & Paddy Nixon

Adaptive Information Cluster, School of Computer Science & Informatics,
University College Dublin, Belfield, Dublin 4, Ireland.

{firstname.lastname}@ucd.ie

ABSTRACT

Group recommender systems introduce a whole set of new challenges for recommender systems research. The notion of generating a set of recommendations that will satisfy a group of users, with potentially competing interests, is challenging in itself. In addition to this we must consider how to record and combine the preferences of many different users as they engage in simultaneous recommendation dialogs. In this paper we introduce a group recommender system that is designed to provide assistance to a group of friends trying to plan a skiing vacation.

Categories and Subject Descriptors: J.7 [Computer Applications]: Consumer Products

General Terms: Human Factors, Design

Keywords: Group recommendation and preference harvesting, Critiquing.

1. INTRODUCTION

In recent years an emerging focus for the recommender system's community has been on improving the effectiveness of these systems through more sophisticated interfacing mechanisms, such as explanation provision, increased transparency, and user modelling [4]. Modern recommender systems engage the user in a conversational dialog with a view to learning about his or her requirements, preferences and willingness to compromise during the course of a recommendation session [3]. Conversational recommender systems of this kind are well-suited to recommendation tasks where a user's initial query is incomplete or vague and user feedback is solicited over a recommendation session in order to develop a more complete view of their requirements [1, 5].

Building a collaborative, conversational recommender that uses the feedback provided by multiple users and which must generate suggestions that satisfy the group introduces two key challenges of interest here: (1) How can multi-user interaction be best supported in a way that facilitates the harvesting of feedback and preferences from multiple simultaneous users?, (2) How best can we model individual and group preferences with a view to influencing recommendation to ensure that the resulting suggestions are likely to

satisfy the individual and the group?

In this paper we describe a novel conversational, collaborative group recommender system called CATS (Collaborative Advisory Travel System), designed to help a group of friends to plan and arrange their skiing vacation. Although CATS will work over an arrangement of PC's, in this paper we describe the system as designed around the DiamondTouch interactive tabletop [2]. Group user feedback is used to suggest products that satisfy the individual and the group. In addition group recommendations are generated proactively through the shared interaction space.

2. GROUP RECOMMENDATION

Our approach to group recommendation is based on a collaborative recommender framework that, at the interface layer, assumes the availability of individual and group interaction spaces and at the recommendation layer, assumes a recommendation engine that is able to record and manage personal as well as group profiles. The group profile is the combination of individual personal preferences and recommendations are made for an individual based on their personal preferences and the preferences of the group. At the same time recommendations can be proactively made for the group with reference to the group preferences.

In this paper, we propose to use the DiamondTouch tabletop device to showcase our group recommender system, CATS (Collaborative Advisory Travel System). CATS will help groups of users find appropriate ski holidays. Our ski holiday dataset consists of 5738 vacation cases. Each case is described in terms of 43 features; of which 25 belong to the resort description (i.e., *beginner*, *cannons*, *transfer time*, *drag lifts*, etc.) and 18 belong to the hotel description (i.e., *stars*, *price*, *balcony room*, *ensuite bath*, etc). In this section we provide an overview of the CATS system, and expand on the interaction and recommendation processes.

2.1 Interaction Component

Critiquing is a form of minimal feedback which helps conversational recommender systems to narrow the search space and help the user find the product they are looking for more efficiently [1, 5]. A critique is a directional feature preference indicated by the user in relation to a presented recommendation. For example, a user of a holiday recommender may specify that they are looking for a similar cheaper holiday by critiquing the *price* feature. The critiquing component of our system is shown in Figure 1(c).

Apart from the simultaneous collaborative aspects of the whole group interfacing with the system on a single Dia-

*This material is based on works supported by Science Foundation Ireland under Grant No. 03/IN.3/I361.

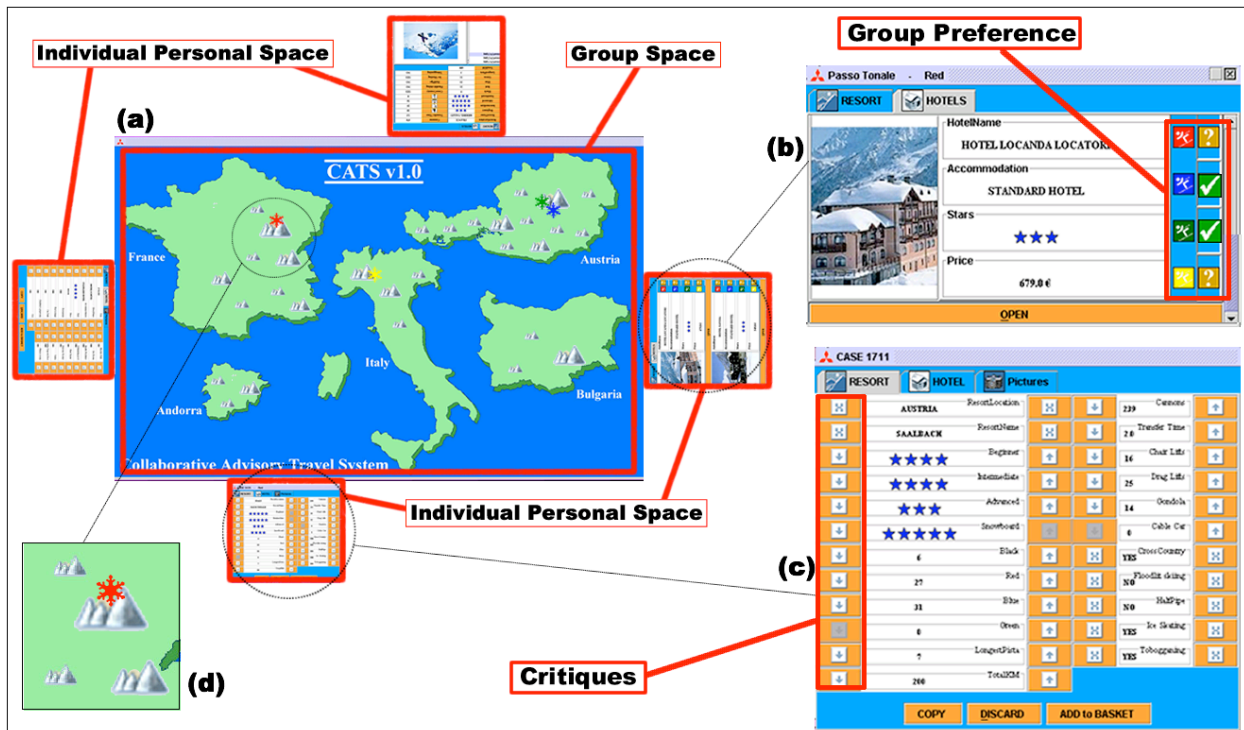


Figure 1: Overview of the Collaborative Advisory Travel System (CATS).

mondTouch, this situation also allows other interesting multi-user interaction. The users are all positioned around the DiamondTouch device. They can copy and pass cases of interest to other users, but they can also confer on a face-to-face basis about their preferences.

Figure 1(a) illustrates the workings of the CATS system. It shows the overall overview of the system, with the group shared map and the individual personal spaces of the 4 users as shown on the DiamondTouch.

2.1.1 Individual Interaction

The individual personal space has some interesting characteristics, that are novel in interfacing terms and that help the group of users, as a whole, to converge on a specific ski holiday decision. In the first of these mechanisms, the hotel overview panel (see Figure 1(b)) lists hotel information depending on the preferred cases of the other users. The basic idea behind ordering the hotel information is to focus the users on the cases where a collaboration or discussion on preferences can be performed by the members of the group. Using the individual personal model of each user, the system counts the number of users interested in that particular hotel and orders the hotels in decreasing order of preferences. At the same time, the hotel overview panel also shows an icon for each user showing their preference for that particular hotel with a check mark or her lack of interest with question mark. With this information the user can see which of the other users in her group have shown a preference for this particular hotel. In addition to these two mechanisms, the user can make a *copy* of the case she is currently critiquing (see Figure 1(c)). The user can then pass the case to another user in the group. Once the

receiving user obtains possession of it, the case is added to their individual personal space. In such a situation the first user is giving awareness of her preferences to another user in order to try and reach a mutual agreement and converge the search for a ski holiday.

2.1.2 Group Interaction

In the group space (the map), there are also some interesting mechanisms that facilitate the group recommendation process (see Figure 1(d)). Firstly, the icon that marks the resort that the user is currently critiquing a case from, has a colour coded snowflake associated with it; each user has a different coloured snowflake. This allows all users of the system to see what resorts other users are currently viewing. In addition, the size of the snowflake changes according to the preferences of the user as stored in the user's individual personal model. This allows all users to judge the *level* of interest of other users in a particular resort. Secondly, the size of the icon that represents each resort grows or shrinks in accordance with the preferences of the whole group.

2.2 Recommendation Component

The system maintains a session-based *group user model* that is made up of those critiques chosen by each user so far and those cases that are still opened. In this section we describe the user models and how they influence both individual and group recommendations. There are two parts to the recommendation component of the system: 1) individual recommendation, where the system reactively recommends cases to the user; and 2) group recommendation, where the system proactively pushes recommendations to the group of users through the group space.

2.2.1 Generating Individual Recommendations

The *group user model* (GUM) is based on user interaction with the system. Each user provides feedback using critiques. These critiques are stored in the $GUM = \{G_1, \dots, G_n\}$, where G_i is a record of a *critique* with its corresponding *user name*. We use the GUM as the basis for our recommendations.

For each recommendation cycle, we break the GUM down into an individual model (IM) of preferences and a remaining members' preferences model (MM). IM belongs to the current user and MM contains the critiques all the remaining members of the group. The individual model is described as $IM = \{I_1, \dots, I_n\}$, where I_i is a single unit critique. The remaining members' preferences (MM) is given by $MM = \{M_1, \dots, M_n\}$, where M_i is a single critique and an associated user identifier. Generating IM and MM, however, is not quite as simple as storing a list of previously selected critiques for the individual or for the remaining members. Users may not always provide consistent feedback, sometimes they make mistakes, change their mind or have incompatible preferences. In this paper, both models are updated in the same way. Specifically, prior to adding a new critique all existing critiques that are inconsistent with it are removed, as are all existing critiques for which the new critique is a refinement [6].

The most important issue for generating individual recommendations is that the system needs to recommend products that contains individual preferences as well as group preferences in a fair way. The key idea is to facilitate future agreement and, at the same time, let the system explore how to maximise preference satisfaction of the whole group.

$$ICompat(c', U) = \frac{\sum_{\forall i} \text{satisfies}(U_i, c')}{|U|} \quad (1)$$

$$GCompat(c', GUM) = \alpha * ICompat(c', IM) + (1 - \alpha) * ICompat(c', MM) \quad (2)$$

$$Qual(c', c, GUM) = \beta * GCompat(c', GUM) + (1 - \beta) * Sim(c', c) \quad (3)$$

In this paper, we propose a user modelling strategy that averages the preferences of the individual and the preferences of the remaining members of the group. The actual recommendation of cases is performed using the incremental critiquing method [6], where a compatibility score for the IM and MM are computed and combined (see Eq. 2) to give a quality score (see Eq. 3). This quality score is then used to rank the candidate case recommendations. This method of recommendation allows us to prioritise those candidate cases that: (1) satisfy the current critique; (2) are similar to the previous recommended case; and (3) satisfy many individual and members previous critiques.

2.2.2 Generating Group Recommendations

In addition to this reactive recommendation of cases, on the basis of explicit user feedback, the group user model (GUM) is also used to bring new cases to the attention of the group in a variety of ways. For example, when a user is viewing the available hotels of a particular resort, the group user model is queried and the hotels are reordered based on critiques contained in the GUM (see Figure 1(b)). In this

way those cases that are most consistent with the preferences of the group are presented to the user in question.

The GUM is also responsible for proactive recommendations that are made via the group space. The objective is to bring potential new sets of cases to the attention of a user, cases that might not be recommended according to the users current critiquing session, but cases that are consistent with group preferences in general. Essentially these new cases are highlighted through the shared interaction space (see Figure 1(d)) by increasing the size of the resort icons associated with those resorts that best match the group preferences learned so far. This occurs each time the group model is updated; that is each time one of the individual users registers a new critique as part of their normal feedback. In this way the central resorts map is continuously being updated to reflect the current group preferences, thereby providing the user with an opportunity to 'break out' of a given recommendation session in order to evaluate a new resort that has been highlighted as a current group preference.

3. CONCLUSION

In this paper, we have introduced CATS, our Collaborative Advisory Travel System, which allows a group of users to simultaneously collaborate on choosing a skiing holiday package which satisfies the group as a whole. This system has been showcased on the DiamondTouch interactive tabletop, which makes it possible to develop a group recommender that can be physically shared between up to 4 users. In the paper we have focused on the core interface, profiling and recommendation issues that have arisen during the development of the system. We have, for example, described how users manipulate their own personal interaction spaces to received personalized recommendations that address their particular needs and emerging group preferences. We have also explained how users interact with a shared space so that they can be alerted to proactive recommendations that are automatically generated by the recommender system based on the developing group model.

4. REFERENCES

- [1] R. Burke, K. Hammond, and B. Young. The FindMe Approach to Assisted Browsing. *Journal of IEEE Expert*, 12(4):32–40, 1997.
- [2] P. Dietz and D. Leigh. Diamondtouch: a multi-user touch technology. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 219–226, New York, NY, USA, 2001. ACM Press.
- [3] D.W. Aha and L.A. Breslow and H. Muñoz-Avila. Conversational Case-Based Reasoning. *Applied Intelligence*, 14:9–32, 2000.
- [4] A. Jameson. User modeling meets usability goals. In L. Ardissone, P. Brna, and A. Mitrovic, editors, *UM2005, User Modeling: Proceedings of the Tenth International Conference*, pages 1–3. Springer, Berlin, 2005.
- [5] K. McCarthy, J. Reilly, L. McGinty, and B. Smyth. Experiments in Dynamic Critiquing. In J. Riedl, A. Jameson, D. Billsus, and T. Lau, editors, *Proceedings of the International Conference on Intelligent User Interfaces (IUI'05)*, pages 175–182. ACM Press, 2005. San Diego, CA., USA.
- [6] J. Reilly, K. McCarthy, L. McGinty, and B. Smyth. Incremental Critiquing. In M. Bramer, F. Coenen, and T. Allen, editors, *Research and Development in Intelligent Systems XXI. Proceedings of AI-2004*, pages 101–114. Springer, 2004. Cambridge, UK.