Adapting Information Delivery to Groups of People

Judy Kay and William Niu

School of Information Technologies
University of Sydney, Australia
{judy, niu}@it.usyd.edu.au

Abstract. When people engage with information, they are often in social groups. This applies, for example, in the case of museum visits, where people typically attend the museum and view the exhibits in small groups. This paper describes a proposed system GM, Group Modeller, which will create group models from a set of individual user models. We discuss the approaches and challenges in terms of common sub-models, collective models, group interaction models, and knowledge-based reasoning across models.

1 Introduction

With computing technology becoming pervasive, information access will become increasingly integrated into normal environments. For personalisation researchers, it is important to begin to take account of the many everyday situations where the user of a personalised service is not alone, but is part of a small group. If we are to customise or adapt information delivery, we can expect that this will operate differently depending on whether a person is alone or in a group. A suitable individual user model may suffice in the situation when a person is alone. It is conceivable that a heterogeneous user group may have conflicting preferences and needs, and the delivery of customised information which addresses the requirements of all the people in the group introduces new challenges in addition to those relevant in the case of an individual. When information is to be adaptive in a group setting, the system will need a group model by appropriately amalgamating the individual models.

There has been some interesting and relevant work on user modelling in the context of personalised museum tours (many of which are reviewed in [1]). As one might expect in this context, some of the work takes account of groups in their design (e.g. [2], [3], and [4]). Petrelli, Angeli, and Convertino [5] observed that people exhibit different behaviour when visiting a museum in a group, as opposed to visiting as an individual. Moreover, a person's behaviour typically depends on the members that make up the group; for example, a young child will typically play a predominant role in determining the duration and the course of a family visit in a museum, with adults accommodating the child's needs and preferences rather than their own.

Another domain that we want to apply the Group Modeller (GM) in is the domestic environment, which differs from museum settings in several aspects (e.g. group composition, session duration, etc.). Unlike a museum, the group members in a domestic environment are usually quite tightly knit groups like families. These people typically have some traits in common (e.g. social values and habits). In addition, and more importantly from a personalisation perspective, since the home is where many people spend large periods of their time, it seems likely that it may be easier to maintain long-term, and perhaps more accurate, user models for people in that context. We are also interested in making use of user models built and maintained in a domestic environment, and then use them in other environments, such as museums.

In the Section 2, we discuss work on group modelling in different domains. Section 3 examines some challenges in composing an effective group model. Next we will describe the proposed group modelling system GM, followed by the conclusions.

2 Group Modelling

Much work has been done in addressing groups in various domains. Masthoff [6] investigated how different group decision rules affect the order of a sequence of numerically rated preferences (in this case, for TV programs) and the satisfaction gained as a whole group by applying each rule. She conducted two experiments. One was, from a third-person view, on how people select a series of TV programs for a group of viewers. In the second one, she examined how satisfied people would feel with the sequences produced by the different strategies. Some interesting findings were that: people tried to account for fairness and avoid individual misery; normalisation was used (i.e. their satisfaction is based on both selected and non-selected items); ratings were judged in a non-linear way (e.g. in a 10-point scale, the difference between 9 and 10 is more significant than that between 6 and 7). This type of work has importance beyond the choice of TV programs. It suggests strategies for dealing with conflicting preferences in more general contexts.

MusicFX [7] is another example of a system that uses individual user models to generate group models. MusicFX is used in a fitness centre to adjust the selection of background music to best suit the preferences of the people exercising at any given time. One interesting facet of the system is that a group in the context is made up of the people who happen to work out at the same time. This is a very different sense of a group from that in most other projects (e.g. [2], [6], [8]), where the group is not composed of strangers; rather it is friends or family members. This system is somewhat unusual in that it uses *explicit* preferences of all participants to make a selection that will directly affect everyone who is present.

Group modelling is also an important factor in some work on city tours. In the INTRIGUE project, for example, attractions are separately ranked by first partitioning a user group into a number of *homogeneous subgroups* with the same characteristics. Then each subgroup may fit one or more stereotypes. Finally, the subgroups are combined to obtain the overall preference, in terms of which attractions to see, for the whole group [9]. In some cases, a subgroup could be particularly influential either because it contains a majority of members of the group or because it represents a relevant tourist class (e.g. children and disabled people). For example, a subgroup may have the following characters: between age of 46 to 55, full mobility, partial vision, and interested in art. One of the stereotypes the subgroup fits is vision-impaired, which would put a high weighting on choosing attractions that facilitate people with impaired vision (e.g. vocal presentation and/or visual aids).

Another well-known group recommender system is the Travel Decision Forum prototype that supports a group of people to plan joint vacations [10]. Inspired by situations where face-to-face communication is not possible, this system emphasises asynchronous group discussions. In the initial phase, each group member specifies his or her preferences by filling in a *preference specification form*. The aim of the next phase is to reach uniformed agreement by having each group member interacting with a virtual mediator, as well as the virtual agents that represent other members. This system introduces a novel type of *incremental preference elicitation*; as each member fills up a preference specification form, he or she may choose to see example solutions based on the preferences of all the group members. Another interesting aspect is how to minimise *manipulative* preference specification. For example, when a person sees the overall rating of an activity is fairly positive, he might rate it lower than what he would have without seeing others' ratings, in order to leverage the final outcome. This still remains a challenging issue involving difficult tradeoffs.

There is rather scant literature on group modelling in the home and museum settings. Sotto Voce [2] was designed to accommodate a group of users in a museum tour, but did not include adaptation as a requirement. Kay, Lum, and Niu [4] presented a scenario on how a scrutably adaptive museum guide may deliver personalised information to each pupil in a school group, which in turn stimulates after-visit group discussions. PEACH is perhaps one of the most ambitious projects in museum research. They discuss issues of how to adapt the information to a small user group. For example, Kruppa [8] explores aspects of providing some common information to the group on a large display and some personalised information on a hand-held display.

Adaptation in homes has involved customisation to the inhabitants as a whole. Voida and Mynatt's [11] experiment on probing families' values reveals a possible approach to designing an information adaptive environment. On the other end of scale, the Casablanca project has designed several prototypes for the home, stressing social communications between family members [12]. Although the devices facilitate communication within the family, it does not provide adaptation for each family member.

3 Challenges in Group Modelling

We describe three of the basic approaches to combining individual user models. Then we describe approaches to enhance reasoning about a collection of individual models in the very likely case that different aspects of people are modelled in the different user models. We need to reconcile these differences to make a more complete and effective group model. We begin by introducing an example. Table 1 lists four hypothetical people's individual user models of preferences on entertainment, which will be referred in the rest of the section. A tick ($\surd$) indicates a positive preference, an X means a negative preference, and a blank space represents an unknown preference. So, for example UM_A models person A, who likes Horror movies, Documentaries, and Musicals, but it does not model the person's preferences for Cartoons or Jazz.

	UM_A	UM_B	UM_C	UM_D
Horror movies	$\surd$		$\surd$	
Documentaries	$\surd$			
Musicals	$\surd$	$\surd$	X	
Cartoons		$\surd$		
Jazz				$\surd$

Table 1. Exemplary Individual Models

Common Sub-models

This is probably the simplest way of combining individual models; it simply involves grouping the properties (or preferences) that all individual models share. We call this a **common sub-model**. For example, when two people, represented by UM_A and UM_B in Table 1 respectively, would like to find a show or movie to watch, this approach would suggest a musical. This approach is essentially a logical AND over the individual models to create the group model.

Collective Models

Unfortunately, a common sub-model often may not provide enough information to compose an adequate group model. In this case, some collective properties have to be chosen to complement the common sub-model. An obvious approach is to perform an operation like a logical OR on the set of individual models to form the group model. We call this process **collective modelling**. Sensibly, the properties that are of the interest to a majority of the group are preferred.

Now take UM_A, UM_B, and UM_C from Table 1 as an example. While there is no one common preference between the three subjects, horror movies and musicals both have two votes. The person with UM_C, however, explicitly expresses a negative preference against Musicals. As a result, a movie with a horror theme may be a more satisfying choice for this group.

Group Interaction Models

The two above approaches are extremely simple-minded. They totally neglect the fact that people react to others in the group. An alternative approach to collective modelling is to account for group composition and social interaction within the group members, which we name **group interaction modelling**. This opens up a wide variety of possibilities and has been a popular research field in group modelling in the past few years. For example, a few group decision strategies discussed in [6], such as the Average Without Misery Strategy, the Fairness Strategy, and the Dictatorship, address group interaction between a group of TV viewers. Those strategies require the group members to have numerical ratings for each property, which represents a TV program in this case.

Another interesting approach that uses the group interaction modelling from a different perspective is the INTRIGUE project [9] reviewed in Section 2. It proposed to partition a user group into a number of *homogeneous subgroups*. Each subgroup had different influential power on the decision making process, and the power might be caused by the size of the subgroup and the class of the subgroup members.

Again, use UM_A, UM_B, and UM_C from Table 1 as an example. While a horror movie may be a best selection using collective modelling, the decision may well be altered if the person represented by UM_B is a 10-year-old child. Because of the adequacy problem, a musical may be chosen despite the subject with UM_C has a negative preference, or either a documentary or a cartoon movie may be chosen under some other group decision strategies.

Knowledge-based Reasoning

The above approaches rely upon the existence of common components in the user models; expressed differently, this means they require that the different user models have a common vocabulary. Where this is not the case, we need inference mechanisms to overcome the problem. This section briefly outlines some of the important forms that this will take.

Figure 1 shows an example of two individual user models. Each circle denotes the user model namespace; it represents all the components modelled. Some of these may not yet be known. For example, UM_1 models user preferences for Tea, Cheese, and Jam, but not Wine. The intersection of the two circles represents the group model between the two individual models. A black dot is a positive preference of a property (e.g. Tea in UM_1). A white dot is a negative preference of a property (e.g. Jam in UM_2). In the case where a component is at one user model but not the other (e.g. Wine in UM_2), it means the component is only modelled in that user model.

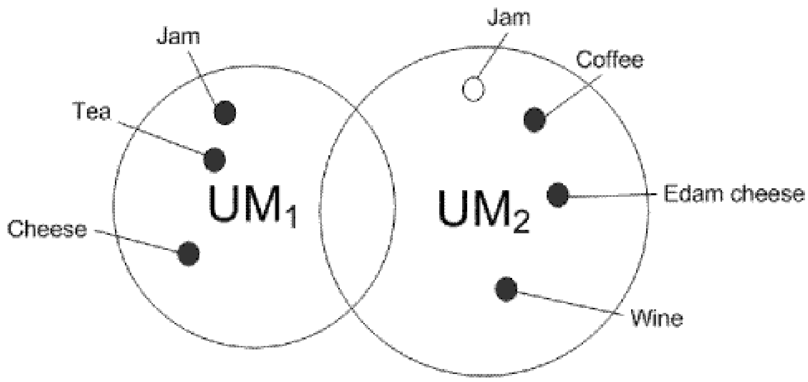

Fig. 1. Exemplary User Models

Common sense reasoning– This is an important part of making sensible assessments about people. We have already alluded to one example of this, in the case of a group which includes adults and children choosing a TV show. It is common sense that the adults will ensure that the children's needs are given priority and adults would expect to watch child-appropriate programs.

Figure 2 gives an example of how the graph would look after applying common sense reasoning to Figure 1. If person 1, represented by UM_1, is a child, he or she normally is not allowed to drink wine, regardless whether he or she likes it or not. So the preference for wine is negative when those two people spend time together.

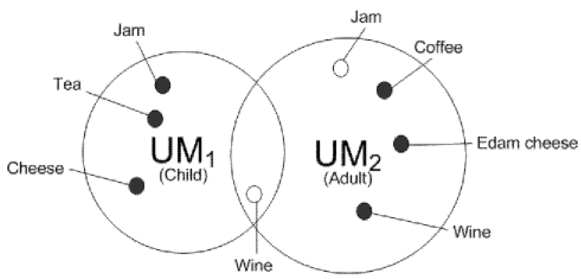

Fig. 2. Applying Common Sense Reasoning

Stereotypic reasoning – Stereotypic reasoning is an important form of user modelling inference that uses statistically valid generalisations to quickly start up a user model [13]. It uses a *trigger*, in the form of simple, readily available information to make a large number of low quality default inferences; they should be overridden once more reliable evidence is available. To give an example of stereotype, if a person is known to be a university professor, it may be suggested that he or she is intellectual, well-educated, fairly wealthy, honest, male, over forty, and well travelled.

This reasoning may also be applied on groups to explore more similarities among a group of people. Say, a group of three people, represented by UM_A, UM_B, and UM_D from Table 1, would like to settle on something to do during a weekend. Note that there is little information about the person with model UM_D. Suppose, we have stereotypic knowledge that people who like jazz typically like musicals and horror movies. In this case, a musical may be recommended as the social activity between the three people.

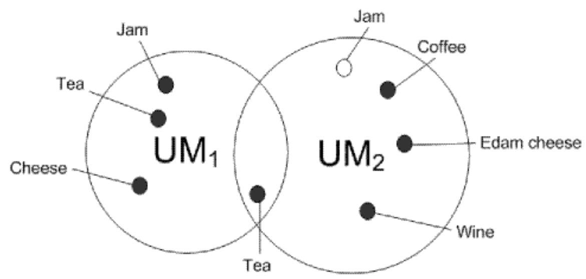

Fig. 3. Applying Stereotypic Reasoning

As a graphical example, Figure 3 illustrates how the example shown in Figure 1 would be after applying stereotypic reasoning. Suppose one of the stereotypes infers that people who like coffee also like tea. This stereotype infers that person 2, who likes coffee, also likes tea. Hence the positive preference of tea for the group model.

Ontological reasoning – Another important form of reasoning for user modelling is ontologically-based. This is important for determining relationships between vocabularies, hence establishing connections across the user models of different individuals. Figure 4 exemplifies the use of an ontology to reason the user models in Figure 1. Person 1 likes all cheese, and person 2 likes the Edam cheese. Hence it is likely that person 1 will like Edam cheese.

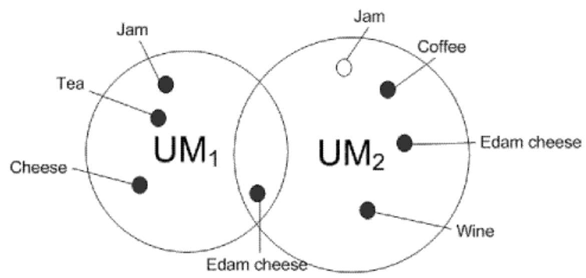

Fig. 4. Applying Ontological Reasoning

4 Proposed System Architecture

Figure 5 illustrates the envisioned architecture for the Group Modeller. We now describe the approaches and tools we propose to use in order to realise it.

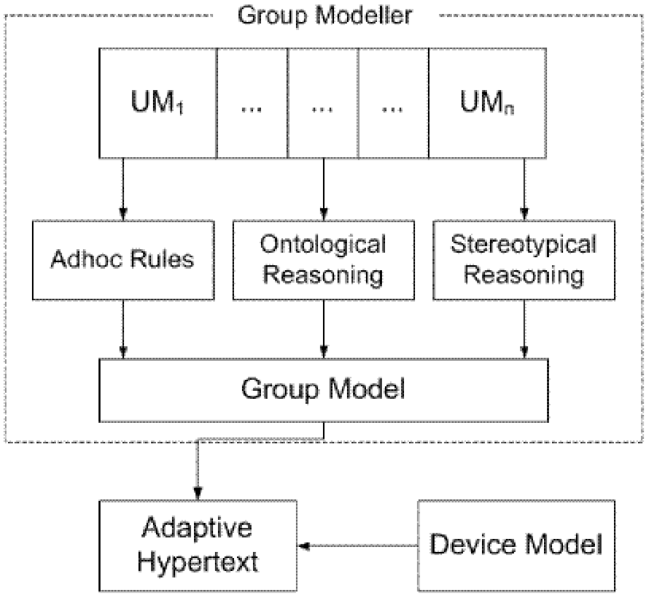

Fig. 5. Envisioned System Architecture

User Modelling Server

The personalisation of information delivery is powered by the user model. Actions by the user are stored in the user model as evidence. The evidence may range from the duration of time the user has spent viewing or interacting with a particular museum exhibit to the history of exhibits he or she has visited. Based on this evidence the system can draw conclusions about user preferences and from this, tailor the delivery of information.

The user modelling server Personis [14] allows adaptive systems to easily manage evidence for user models, and provide a resolution system to conclude a value of each user model component based on this evidence. These resolvers are crafted by the system designers with scrutability in mind. At any time, users should be able to ask the system why an adaptation was performed, and the system should respond with the evidence that lead to the adaptation. With this in mind, the same resolvers can be accessed by different devices, with the results tailored at the device level to be appropriate to the interface. These properties also make Personis a suitable candidate for modelling groups. We simply need to establish suitable approaches for the new resolvers that will reason about groups, rather than just individuals.

Ontological Reasoning

We propose to model this knowledge with a light-weight ontology. MECUREO [15] is one such tool to fit this task. It was originally designed to create an ontology of computer science terms from the Free On-Line Dictionary Of Computing (FOLDOC) and has since been used in several experimental systems, using various dictionaries/glossaries.

Stereotypic Reasoning

In the Personis approach, this is managed by a knowledge source that provides evidence. That evidence is distinguished as stereotypic, and resolvers treat this as less reliable than other forms of evidence. We envisage that the Group Modeller would need to establish which components in the various individual user models need additional evidence. Backward chaining through stereotypes could be used to search for suitable stereotypes to support reasoning about these for the group.

Adaptive Hypertext

We have been developing a version of the Scrutable Adaptive Hypertext system [16] that integrates the Personis user modelling server. The web-based interface, adaptability, and controls for scrutability make it a suitable medium for the system described above. Each page is tailored to the user(s). Whole pages may be omitted. At any time the user may choose to see how the page currently viewed is adapted to her or him. The text being included or excluded is highlighted with different colours. By moving the mouse cursor over each section of the highlighted text, the reason for inclusion or exclusion is provided. A basic description of the user model is also displayed to the right of each page.

5 Conclusions

There are many situations where information should be delivered to a group of users. This requires that we manage group models. We have explored some of the issues involved in doing this, identified some approaches that should be part of a Group Modeller and presented the architecture of an experimental system that will be a testbed for group modelling.

Acknowledgements

This project is partially funded by Smart Internet Technology CRC.

References

1. Filippini-Fantoni, S.: Museums with a personal touch. In: EVA 2003 London Conference Proceedings. (2003)
2. Aoki, P.M., Grinter, R.E., Hurst, A., Szymanski, M.H., Thornton, J.D., Woodruff, A.: Sotto voce: Exploring the interplay of conversation and mobile audio spaces. In: CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems, Minneapolis, Minnesota, USA, ACM Press (2002)
3. Stock, O., Rocchi, C., Zancanaro, M., Kuflik, T.: Discussing groups in a mobile technology environment. In A. Buts, C. Kray, A.K.A.S.H.P., ed.: MU3I '05: 2nd Workshop on Multi-User and Ubiquitous User Interfaces 2005 in Conjunction with the IUI'05, San Diego, California, USA, ACM Press (2005)
4. Kay, J., Lum, A., Niu, W.: A scrutable museum tour guide system. In Butz, A., Kray, C., Kruger, A., eds.: MU3I'05: 2nd Workshop on Multi-User and Ubiquitous User Interfaces at IUI '05. (2005)
5. Petrelli, D., Angeli, A.D., Convertino, G.: A user centered approach to user modelling. In Kay, J., ed.: UM'99: Proceedings of the Seventh International Conference on User Modelling, Springer Wien New York (1999)
6. Masthoff, J.: Group modeling: Selecting a sequence of television items to suit a group of viewers. User Modeling and User-Adapted Interaction **14** (2004)
7. McCarthy, J.F., Anagnost, T.D.: Musicfx: an arbiter of group preferences for computer supported collaborative workouts. In: CSCW '98: Proceedings of the 1998 ACM conference on Computer supported cooperative work, Seattle, Washington, United States, ACM Press (1998)
8. Kruppa, M.: A user study on parallel presentations for heterogeneous user groups. Technical Report 563, Saarland University
9. Ardissono, L., Goy, A., Petrone, G., Segnan, M., Torasso, P.: Intrigue: personalized recommendation of tourist attractions for desktop and hand held devices. Applied Artificial Intelligence **17** (2003)
10. Jameson, A.: More than the sum of its members: challenges for group recommender systems. In: AVI '04: Proceedings of the working conference on Advanced Visual Interfaces, Gallipoli, Italy, ACM Press (2004)
11. Voida, A., Mynatt, E.D.: Conveying user values between families and designers. In: CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems, Portland, Oregon, ACM
12. Hindus, D., Mainwaring, S.D., Leduc, N., Hagström, A.E., Bayley, O.: Casablanca: designing social communication devices for the home. In: CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press (2001)
13. Kay, J.: Lies, damned lies and stereostypes: pragmatic approximations of users. In Kobsa, A., Litman, D., eds.: UM '94: Proceedings of the Fourth International Conference on User Modeling. (1994)
14. Kay, J., Kummerfeld, B., Lauder, P.: Personis: A server for user models. In: AH'02: Proceedings of Adaptive Hypermedia and Adaptive Web-Based Systems, Springer-Verlag, London, UK (2002)
15. Apted, T., Kay, J.: Mecureo ontology and modelling tools. WBES of the International Journal of Continuing Engineering Education and Lifelong Learning **14** (2004)
16. Czarkowski, M., Kay, J.: A scrutable adaptive hypertext. In: AH '02: Proceedings of the Second International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, Springer-Verlag (2002)