Using a Domain Ontology to Mediate between a User Model and Domain Applications

Charles Callaway and Tsvi Kuflik

[1] University of Edinburgh, 2 Buccleuch Place
Edinburgh, Scotland EH9 1SA
ccallawa@inf.ed.ac.uk
http://homepages.inf.ed.ac.uk/ccallawa/

[2] MIS department, University of Haifa, Mount Carmel, Haifa 31905, Israel
tsvikak@is.haifa.ac.il
http://mis.hevra.haifa.ac.il/~tsvikak/Home.htm

Abstract. User modeling data in dynamic, personal and adaptable systems is usually collected immediately before system interaction with a questionnaire, or during application execution when users' choices are recorded and analyzed. This data is then typically used to intelligently adapt the system's output, hopefully improving the user's interaction in some measurable way. When coupled with knowledge-based applications such as intelligent information presentations or tutoring systems, this user model may be mapped onto the system's knowledge base as an *overlay* that may describe what domain material has been experienced by the user or to adaptively encode a progression of topics to be presented next. We present a case study in the museum domain, where an adaptive hypermedia mobile presentation system creates a user model for its own use, and subsequently a post-visit report generation system modifies the data in the model to produce a personalized summary of the entire museum visit. We describe how one component, the interest model, is seeded by the knowledge-based user modeling data collected in the initial mobile phase, and is then expanded via inference over the existing domain ontology during the second phase of report generation.

1 Introduction

The need for information presentation systems to automatically adapt themselves to their users is recognized in many application areas, among them visitor guides for tourists like CyberGuide [9] or Deep Map [6], and museum visitor guides as developed for projects such as Hips [12]. In order to adapt a system to individual users, there is a need to identify each user's needs and to model the user in order to guide the adaptation process. Such a user model should provide the information needed by the specific application, and hence usually contains only relevant application-related data.

Information about users can be collected in two ways: explicitly, by asking the user specific questions in order to provide the relevant information (an approach adapted by the GUIDE context aware tourist electronic guide which asked its users to provide some personal details such as their name, interests and preferred reading language [5]) or implicitly, by monitoring users' behavior and inferring a user model based on a number of observations, an approach taken for example by researchers for the HIPS system in the non-intrusive version of their museum visitor's guide [10]. A major drawback of explicit user modeling is the need for users' active involvement, an effort users typically seek to avoid given the need for asking a number of potentially intrusive, private questions. Implicit approaches are also troublesome given that inferring user interest and knowledge from their behavior is a highly uncertain task.

Using stereotypes – a representation of clustered groups of users – is one way to partially overcome these limitations: users are requested to provide a small amount of personal (but not identifying) information that allows the system to assign them to a stereotypical group from which relevant user details can be inherited. Such an approach was taken in the INTRIGUE project [2], a mobile tourist guide for the city of Turin. INTRIGUE recommended destinations and itineraries for family groups, allowing for various points of view such as historical period and artistic themes. It elicited user desires with direct questions, storing demographic and background data for each person in the group, and then using a probabilistic user model to predict their joint interests.

Like other guide domains, the museum environment [17] is a challenging environment for user modeling. First, extensive and detailed information about the museum exhibits is needed in order to provide the relevant knowledge space to match the wide array of possible presentations generated for particular users on their individual paths through the physical space. In addition, there is very little information (if at all) about the visitors entering the museum, and it is impossible in a real museum environment to ask them to explicitly provide all the information required to create an extensive user model. Hence non-intrusive user modeling that can still yield adaptive, personalized results is required.

Specific challenges include user modeling, as described by Serini and Straparava [19] in the HIPPIE project, and has been the subject of a wide range of research efforts with well-developed user models ranging from kiosk-based clients [1] to mobile PDAs [18]. HIPPIE represented a user by a user model that took into account the user's personal data (such as age, job and more), his level of domain knowledge and his interests. The unique setting and the complex nature of the museum together with the highly varied characteristics of the visitors led researchers to suggest very personal support would be preferable to the stereotype based approach, which was confirmed in subsequent highly detailed but unimplemented research on the HYPERAUDIO system [13].

A significant advantage of the deep representational level inherent in knowledge bases is the availability of a taxonomy of concepts (ontology) for use in semantic processing, for instance to gauge similarity between two domain concepts or to record exactly what information has already been presented to the user. This is true in other domains needing user modeling beyond museum or tour guides such as intelligent tutoring systems [11] and e-learning [7].

We describe two of our museum applications in the PEACH project which both share a common knowledge base, but employ a user model in different ways. Our KB supports typical functions such as reasoning for overlays and for determining sets of subsequent potential choices for the user to explore next. In addition, it also supports a sophisticated natural language generation interface which immediately after a museum tour allows us to create printed or emailed reports for museum visitors, consisting of a record of their individual interests and recommendations for followup learning, which they can then read at home or online after their museum visit. The report generator takes the original user model created implicitly during the visitor's tour, makes extensive use of the existing KB ontology to infer a personalized interest model, and creates the report detailing what he or she might like to see during their next visit to that museum or other similar nearby museums.

2 The PEACH Domain, Mobile Guide, and Report Generator

A visit to a museum in the context of the PEACH project consists of two separate phases: an immersive visit through the museum accompanied by a PDA which transmits interface events and receives animated, interactive presentations for the visitor [16], and a natural language report generator [3] which delivers a personalized electronic or paper report describing the visitor's trip through the museum along with additional information that the accumulated interest model predicts that that particular visitor would like to see. The user model is formed by the individual path the visitor takes through the museum and the choices for information on new objects they made via the PDA interface [8].

As an application domain we have chosen the *Cycle of the Months* of the Torre Aquila at the Buonconsiglio Castle in the city of Trento, Italy. This work is composed of eleven side-by-side frescos (measuring on average 2 meters wide and 3 meters high each, and representing a particular calendar month, such as January in Figure 1) painted during the 1400s and illustrates in great detail and complexity the activities of medieval aristocrats and peasants of Trento throughout a full year. Each fresco presents several scenes that were typical of daily life in that month, such as aristocratic, plebeian, and religious activities. The frescos are highly detailed, depicting architectural details, clothing and tools of the era, providing an environment rich in similarities and differences which can be exploited by the user model and domain knowledge base.

The underlying knowledge base contains around 1500 domain concepts for each visible object in a fresco (people, animals, buildings, etc.), properties of those objects (color, size, history, spatial position, etc.), active and stative relations (jousting, building, sitting, etc.), plus a larger number of generic concepts necessary for generating explanations with the report generator. Additionally, these concepts are organized under an *ontology* that separates concepts of unlike types and allows an inference engine to determine which concepts are similar, and thus might be of interest to visitors. The ontology is coded directly into the

Fig. 1. The January fresco of the Torre Aquila in Trento, Italy

KB as taxonomic relations between concepts of varying abstraction. The KB was built completely by hand, requiring around three months of effort.

2.1 The Mobile Guide

The interactive museum tour takes place with an adaptive mobile museum visitors guide developed within the PEACH project. The electronic guide provides museum visitors with dynamically generated personalized presentations relevant to the exhibits they see in the Torre Aquila. They include details about specific scenes found in the frescos, presented from a certain perspective (such as artistic or historical). The presentations are based on detailed information that is drawn from the system's knowledge base, and is selected based on user preferences as inferred by the system [15]. Presentations are given to the visitors on a choice of devices - a PDA that the visitor carries, or a desktop with better displaying capabilities, if the visitor is near such a device when a relevant presentation is prepared. The personalization is based on a set of features representing information continuously gathered about the visitor during the museum visit, including:

- *Spatial information*: whether the user has already been in this area before, is facing an artwork, or she/he is spatially nearby some exhibit;
- *Interests*: whether the user is interested in a specific artifact;

– *Discourse history*: what particular shots the user has already seen, or if she/he has already watched a presentation for an artwork;
– *Device*: if the user is visiting a museum provided with a PDA, or she/he is looking for information on the web with a desktop PC;

The PEACH dynamic user modeler works in a "non-intrusive" manner. Hence there is no information about the visitor when starting each visit, and as a result, the model is built solely by observing the visitor's behavior. From the beginning, the user is tracked by recording their positions (in terms of the visited exhibits), the time spent at each position, and the details about the presentation delivered (the main theme, the global perspective and peripheral concepts included). User interests are defined in terms of domain concepts, which are associated with individual presentations. These concepts provide a description of the content of the presentation, thus representing its theme. The concepts that are associated with the presentations and used for modeling user interests are drawn from a domain knowledge base that is primarily designed for natural language generation for visit summary reports. Since there is no prior knowledge about the user, the knowledge base and the concepts associated with the individual presentation are the only source of information for user preferences with respect to the exhibits visited and presentations delivered in the current museum visit.

The information described above is continuously being gathered and the level of user interest inferred by the user model working in a "non-intrusive" manner. User level of interest is inferred by monitoring each user's explicit and implicit feedback. Explicit user feedback is given by pressing a "More" button (to indicate a positive reaction) or an "Enough" button (for negative reaction) and implicit feedback, in the form of presentation completed without interruption (positive reaction). These are used to infer user interest in the various concepts presented by the presentations that the visitor sees. Using an inference mechanism that follows ontological links in the knowledge base, user interests are propagated from the specific concepts associated with the presentations to more abstract, related concepts (*e.g.*, interest in a concept "knight" is propagated to the more abstract concept "aristocracy").

Explicit feedback has a higher priority than implicit feedback in the sense that explicit feedback is more reliable so it drives an immediate change in level of interest in the concepts associated with the delivered presentation, while implicit feedback requires accumulation of evidence for every concept (several implicit responses) before changing a visitor's interest level in that given concept. In addition to the level of interest, a "certainty factor" is used as a way of representing the semantic distance between the original concept where the inference started and the current, inferred concept. Concepts which are part of presentations are seeded with an initial neutral value – "interested a little". Later interactions change the model of the visitor's interests in the various concepts using a qualitative 5-level scale (for more details, see [8]).

2.2 The Report Generator

Supporting adaptive, intelligent presentation generation requires recording events during the visit, updating information about user interests, and choosing new presentations based on these inferred user interests, as described above. Generating a subsequent report about that visit instead requires a different perspective on the type of information presented and the types of inferences used because we want to (1) explicitly describe in detail exhibits that seemed to be of special interest to the visitor, (2) compare and contrast details in order to increase their understanding of what they have seen, (3) describe museum exhibits the visitor didn't see in addition to those seen (in order to trigger interest for future visits), and (4) recommend a number of exhibits both in this and in other museums so as to enrich the overall user experience.

To create the text of the report, the report generator combines a text planner that determines the most relevant information to put into the description along with its coherent rhetorical organization, and a deep syntactic NLG system that creates the actual text read by the visitor. The text planner accesses the user model, ensuring that the resulting text will be personalized, and makes extensive use of the knowledge-based ontology and ontological inference mechanism described below to decide what were the favorite exhibits seen by the visitor as well as what exhibits might have been favorites if the visitor had had enough time to see them.

For adaptive generation that is highly personalized for a particular museum-goer's visit, it is important to ensure a high amount of variation in the resulting text. To achieve such variability, the text planner queries the user model to get the log of the user interactions. For instance, to *sequentially* describe what the visitor saw, the text planner extracts an ordered list of visited artworks and accesses the knowledge base to get a shallow description of the main contents of each artwork to be included in the summary.

Alternatively, the text planner can retrieve a list of ranked topics from the inferred interest model. Thus the corresponding *thematic* report may consist of a series of paragraphs describing the top items in the interest model. To prevent excess repetition of similar exhibit types (for instance, where the entire text is about the knights, lords, and ladies in the frescos), we cluster semantically related items in the interest model, making extensive use of the ontology to inform us which items are conceptually related to others. The text planner also includes additional details from the knowledge base based on perspectives the user seemed globally interested in. For example, if the visitor spent a lot of time watching and requesting information on castles, churches and buildings in the frescos, the ontologically informed heuristic assumes the user is interested in architecture and thus includes architectural facts from the knowledge base, such as the name of the builder or its particular building style.

The text of the report is created using the language-independent StoryBook deep NLG system [4], which handles low-level language issues, such as sentence subjects, verbs, pronouns, morphology, etc. Deep NLG has advantages that make it useful in generating extended reports in a museum context: the text can be

Visit Log / User Model	Initial Interest Model	Extended Interest Model
Moved in front of January for 223.0s Started January-Fresco Overview Completed Presentation on HUNTERS Completed Presentation on CASTLE Stopped Presentation on CASTLE-WALLS	HUNTERS + HUNTING-DOGS + BADGERS + SNOW + CASTLE + CASTLE-WALLS -	ARISTOCRACY + ARIST.-ACTIVITIES + ANIMALS + WINTER + ARCHITECTURE - BATTLE -
Moved in front of February for 192.0s Started February-Fresco Overview Completed Presentation on TOURNAMENT Completed Presentation on KNIGHTS Stopped Presentation on SPECTATORS Completed Presentation on BLACKSMITH	TOURNAMENT + KNIGHTS + LANCES + HORSES + SQUIRES + SPECTATORS -	ARIST.-ACTIVITIES + ARISTOCRACY + WEAPONS + ANIMALS + SOCIAL-ROLES + ARISTOCRACY -

Sequential, Non-Adaptive Visit Paragraph	Thematic, Adaptive Visit Paragraph
... You first went to see the January fresco which contains many scenes of winter activities. The main theme of this fresco is a snowball fight between a group of nobles in the bottom panel. Two hunters are leading their dogs to search for badgers, while a lord is cutting roses in his castle garden ...	... Your favorite theme was the activities that the nobles engaged in during their daily lives. For instance, the hunt in the snow in January, the knights engaged in the tournament in February, and September's hawk hunts captured your attention for a large part of your visit ...

Table 1. Inferring an interest model from a visit log.

in multiple languages, produced in high-quality prose, provide for automatic variation at the syntactic and lexical levels, and contain integrated markup (such as HTML, or prosody for text-to-speech). Other NLG systems have had user modeling components, such as the STOP [14] report generation system; though our implementation is meant to work in multiple languages and does not require any explicit information from the user. Importantly for the production of printed, color reports that will be read by museum visitors, the NLG system allows the introduction of HTML markup into the text at the syntactic level and thus can produce the report as a web page which can also be emailed [3], and the report includes an image of each artwork the visitor was interested in.

3 Inferring the Interest Model using a Taxonomy for Presentation Generation and Report Generation

Between the completion of the interactive tour and the creation and printing of the report as the visitor leaves, the user model that results from the tour must be extended to allow for inference of the interest model that covers more than the exhibits seen by the visitor. Otherwise, the visitor would receive a report that is merely a copy of the visit log, describing the sequence of exhibits visited.

For instance, imagine a visitor who carries their PDA in front of the January fresco, as described in Table 1, and watches a series of presentations about aristocrats enjoying leisurely winter activities (Figure 1). The visitor watches entire presentations about these topics, but interrupts the PDA when it begins to describe the architectural details in the same scene. The visitor then moves on to the February fresco, and again listens patiently to presentations about aristocratic activities in addition to those of the lower social classes. The visit log collects this data from the mobile presentation system, and integrates it with the knowledge base, producing an initial interest model.

From the joint knowledge that the visitor enjoyed a hunting scene in the January fresco and a tournament in the February fresco, we can generalize to infer that the user is interested in winter aristocratic leisure activities. Integral to this inference is ontological knowledge as well as other types of relations that connect concepts into a large semantic net. The knowledge base records information about each scene being presented, for instance that the presentation on hunters includes the hunting dogs, a badger, the hunters themselves, etc. Later, after inference, the text planner may additionally cluster multiple similar conceptual interests to produce abstract topics such as "aristocrats" rather than "knights" and "ladies" individually.

If well-defined clusters can be created for the top-rated concepts at the completion of inference for the interest model, the text planner then writes a report organized thematically and centered on the top clusters in the list [3]. Otherwise, the system chooses a sequential report describing exhibits in order, pulling details from the interest model. In our experience, a thematic organization of the report is superior to sequential methods of organization. In either case, adaptive variation in the report text is ensured at the organizational level, depending on the visitor's path through the museum and their requests for further details. Finding clusters is also important for report generation because without them the final report may consist of redundant sequences of text talking about similar interests without providing any contrasting or unifying information.

4 Discussion

In our application, a single user model supports two different tasks: online presentation generation and creating personalized summary reports. On one hand, the user model supporting adaptive presentation generation requires recording events during the visit and abstract information about user interests, which are more general by nature than the specific concepts depicted by the presentation. For instance, the specific jousting activity of knights may be a concept associated with a presentation, but for user modeling purpose the more abstract concept of aristocratic activity is what really matters. On the other hand, the report generator requires much more detailed information than is available in such interest models; the required knowledge should include details such as the number of jousting knights in the scene, their weapons, clothes, relative position to each other and so on in order to support potentially all possible reports from every visitor perspective. To provide this level of detail for an application, we must overcome a gap between the standard functions a user model supports for abstract user modeling and the domain-specific expectations of the application.

This gap can be seen in what was initially available to report generation: an unordered or semi-ordered list of interests extracted from the knowledge base with associated discrete interest annotations. The list included the specific concepts and related, more abstract concepts. For dynamic presentation generation, all that is needed is a list of concepts and a level of user interest in each and every one of them, so that the user model can respond to queries such as what is the level of interest the user has in some concept X (in addition to queries like whether the user has seen a concept X or did the visitor visit some exhibit X). Such a list is ill-suited for direct processing by a report generator, which needs rhetorical and discourse motivations to produce text. Given the list-like nature of the interest model, driving the text planner thus required a number of basic list functions whose parameters include the semi-ordered topic list, the knowledge base, and fundamental user parameters. However, standard implemented user models do not support these functions:

- *Filtering*: Removing or retaining particular interests that satisfy a filter condition, such as all artwork elements containing animals or farm implements.
- *Clustering*: Grouping similar interests under more abstract hierarchies (which may need to be constructed on the fly), such as aristocrats from a set of interests including lords, ladies, and knights, to avoid repetition in the report.
- *Sorting*: Placing a series of interests in some logical order, so that the report doesn't result in a sense of random order.
- *Splitting*: Separating similar items into groups depending on an external element, such as when they are distributed across adjacent artworks, and especially for incorporating rhetorical effects like comparison and contrast.
- *Searching*: Looking through the knowledge base for items similar to a given interest, which can be used to populate the text with additional details.

Such services are required for the system to group together the individual concepts in a semantically meaningful way. Such a grouping lets the system focus on concepts that were the most interesting for the visitor and for later elaborating specific concepts of interest by querying the knowledge base for more details.

5 Conclusions

Specific implementations impose specific requirements on user modeling. In this paper, we described a method to bridge the gap between an abstract user model needed for a dynamic museum guide system and a detail-centered interest model needed for a report generator. The domain knowledge base served as a foundation for this bridge, allowing implicit visitor behavior to determine an adaptive, personalized report of the visitor's experience.

References

1. Calder, J., A.C. Melengoglou, C. Callaway, E. Not, F. Pianesi, I. Androutsopoulos, C.D. Spyropoulos, G. Xydas, G. Kouroupetroglou and M. Roussou: Multilingual Personalized Information Objects. Multimodal Intelligent Information Presentation. O. Stock and M. Zancanaro (eds.), Series: Text, Speech and Language Technology, Vol. 27, Kluwer Academic Publishers, 2005.

2. Ardissono, L., A. Goy, G. Petrone, M. Segnan and P. Torasso: INTRIGUE: Personalized Recommendation of Tourist Attractions for Desktop and Handset Devices. *Applied Artificial Intelligence* 17(**8-9**):687–714, 2003.
3. Callaway, C., T. Kuflik, E. Not, A. Novello, O. Stock and M. Zancanaro: Personal Reporting of a Museum Visit as an Entry-point to Future Cultural Experience. IUI 2005.
4. Callaway, Charles and James Lester: Narrative Prose Generation. *Artificial Intelligence* 139(**2**):213–252, 2002
5. Cheverst, K., N. Davies, K. Mitchell, A. Friday and C. Efstratiou. Developing a Context-aware Electronic Tourist Guide: Some Issues and Experiences. Proceedings of the SIGCHI conference on Human factors in computing systems, The Hague, The Netherlands, pp. 17–24, 2000.
6. Fink, J. and A. Kobsa: User Modeling in Personalized City Tours. *Artificial Intelligence Review* 18(**1**):33–74, 2002.
7. Kay, J., and A. Lum: Ontologies for Scrutable Student Modelling in Adaptive E-Learning. Proceedings of the Adaptive Hypermedia and Adaptive Web-Based Systems Workshop on Semantic Web for E-Learning, August 2004.
8. Kuflik, T., C. Callaway, O. Stock and M. Zancanaro: Non-Intrusive User Modeling for a Multimedia Museum Visitors Guide System. Submitted to UM 2005.
9. Long, S., R. Kooper, G. Abowd and C. Atkinson: Rapid Prototyping of Mobile Context-Aware Applications: The Cyberguide Case Study. Proceedings of the 2nd ACM International Conference on Mobile Computing and Networking, 1996.
10. Marti, P., A. Rizzo, L. Petroni, G. Tozzi, and M. Diligenti: Adapting the museum: A non-intrusive user modeling approach. Proceedings of the Seventh International Conference on User Modeling, pp. 311–314, 1999..
11. Mizoguchi, R. and J. Bourdeau. Using ontological engineering to overcome common AI-ED problems. *International Journal of Artificial Intelligence in Education* 11:107–121, 2000.
12. Not, E., M. Sarini, O. Stock, C. Strapparava and M. Zancanaro: Information Adaptation for Physical Hypernavigation. In Proceedings of the I3 Annual Conference, Nyborg, 1998.
13. Petrelli, D., A. De Angeli and G. Convertino: A User-Centered Approach to User Modelling. Proceedings of the Seventh International Conference on User Modeling, pp. 255–264, 1999.
14. Reiter, E., R. Robertson and L. Osman: Lessons from a Failure: Generating Tailored Smoking Cessation Letters. *Artificial Intelligence* 144:41–58, 2003.
15. Rocchi, C. and M. Zancanaro M: Rhetorical Patterns for Adaptive Video Documentaries. Proceedings of Adaptive Hypermedia and Adaptive Web-Based Systems AH-2004, Eindhoven, 324–327, 2004.
16. Rocchi, C., O. Stock, M. Zancanaro, M. Kruppa and A. Krueger: The Museum Visit: Generating Seamless Personalized Presentations on Multiple Devices. Proceedings of IUI-2004, Madeira, 2004.
17. Stock, Oliviero: Language-Based Interfaces and Their Application for Cultural Tourism. *AI Magazine* 22(**1**):85–97, 2001.
18. Stock, Oliviero and Massimo Zancanaro: Intelligent Interactive Information Presentation for Cultural Tourism. Invited talk at the International Workshop on Natural, Intelligent and Effective Interaction in Multimodal Dialogue Systems. Copenhagen, Denmark. June, 2002.
19. Sarini, Marcello and Carlo Strapparava: Building a User Model for a Museum Exploration and Information-Providing Adaptive System. Proceedings of the 2nd Workshop on Adaptive Hypertext and Hypermedia, Pittsburgh, PA, June, 1998.