

Adaptively Recommending Museum Tours

Adrian Bright, Judy Kay*, Daren Ler*, Kelvin Ngo, William Niu*, and Alfonse Nuguid

School of Information Technologies
University of Sydney
NSW 2006, Australia

*{kay, ler, niu}@it.usyd.edu.au

ABSTRACT

We believe that a user should be able to *scrutinise* the beliefs a system has about him/her, namely the user model, as well as be able to update that model. A MyMuseum prototype of a scrutable and personalised museum guide has been built to explore ways to visualise how the system adapts to the user and the reasons for the adaptation. As a result of the evaluation, we propose to improve the process of recommending a touring path by using an automatically built dynamic ontology.

Keywords

User modelling, museum, scrutable system, multi-ontology, recommendation

1. INTRODUCTION & BACKGROUND

According to Petrelli [7], most museum visitors welcome some sort of guidance. That is, to some extent they would prefer to receive recommendations regarding what to see, instead of browsing the museums autonomously. Despite this, most of the electronic museum guide prototypes that have been developed so far do not include personalised recommendations concerning the sequence of exhibit(s) to visit. Most of them instead focus on audio-visual augmentation (e.g. [2], [4]), unobtrusive personalisation of information delivery (e.g. [6], [10]), or less user-oriented technicalities (e.g. [8], [3]).

Museum curators and exhibition designers try to arrange artefacts strategically so as to facilitate coherent tours. However, exhibits may be organised in a number of different ways. For example, an archaeology museum may choose to place artefacts with the same origin in the same vicinity (e.g. a mummy may be placed next to an Egyptian sarcophagus). A visitor, on the other hand, may come to the museum to see the wide variety of ancient ceramics. In this case, the visitor will have to look for the ceramics among the other exhibits and might even miss some of them. Since

a physical museum can have only one physical layout, it is desirable to be able to tailor a touring path according to the visitor's interests and goals of the visit. While it is neither economical nor feasible to provide each visitor with a human guide, some of the flexibility and personalisation of a one-to-one tour would be a valuable way to enhance the museum experience.

While the concept of personalised information delivery is becoming widespread, the idea of being able to *scrutinise* how the information has been adapted is uncommon. We have been exploring ways to build personalised interfaces that put users in control of the personalisation by giving them control of the machine's model of them (i.e. their individual user models). This means the users should have the right to scrutinise their own user models and update them if they wish to. In addition, we want to support user scrutiny over the personalisation process. This would require a way to visualise the user model, as well as the way it is used.

We have built a MyMuseum prototype of a scrutable museum guide. It provides a mechanism for the user to see how information delivered is adapted to him/her and why. The next section gives an overview of MyMuseum, which is followed by the architecture of this system. In Section 4 we present a user study that we have conducted. And finally, we conclude with some insights derived from this evaluation.

2. SYSTEM OVERVIEW

MyMuseum was intended to 1) build a nomadic museum guide that tailors the presentations to the user's interests and preferences, and at the same time 2) allowing the user to visualise what has been adapted to him/her and why. The testbed we used was the Nicholson Museum¹, the first Australian archeological museum, located at the University of Sydney campus.

Before using the tour guide, the user must be authenticated in order for the system to apply the corresponding user model. In the case of a first-time user, the system asks about the user's goal of visit. Then it initialises the user model by making stereotypical assumptions [9] for the value of each user model component. For example, if a first-time user indicates that the goal of her visit is to learn more about Troy, the system will make inferences that the user wants more detailed descriptions for each artefact about Troy, as well as external references/sources for it.

The following elements are made adaptive in this proto-

COPYRIGHT NOTICE: Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Smart Environments and Their Applications to Cultural Heritage, UbiComp 2005, September 11–14, 2005, Tokyo, Japan.

¹The Nicholson Museum: <http://www.usyd.edu.au/nicholson/> (accessed on 23 June 2005).

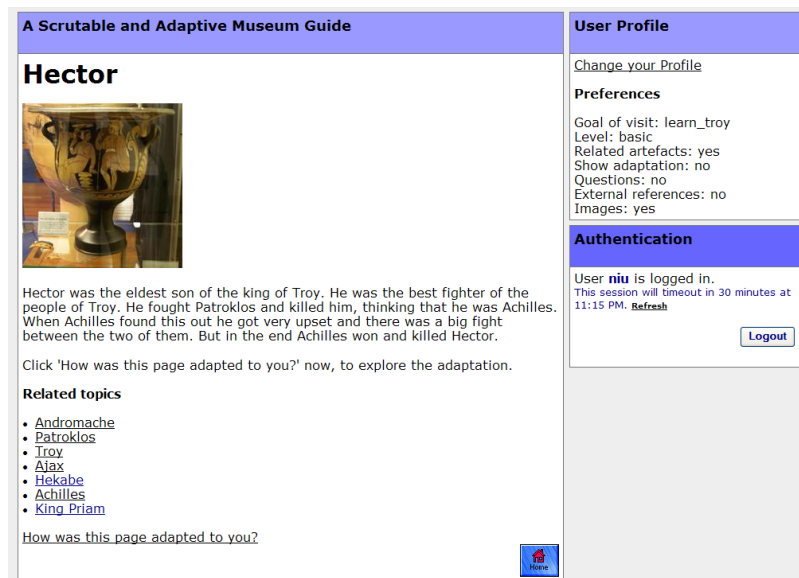


Figure 1: An Example Content Page

type:

- The level of artefact description: basic/advanced.
- Displaying questions for the artefact: yes/no.
- Displaying external links/resources: yes/no.
- Displaying related artefacts: yes/no.
- Displaying image of the artefact: yes/no.
- User's goal of visit: "learn about Troy" / "general museum visit" / none

Figure 1 demonstrates a page of the system that adapts to a user. The main *cell* displays adapted information about an artefact, in this case, a vase about Hector. A name, an image, and a brief description of the artefact, as well as some links to the related artefacts are shown according to the user model. The top right hand cell with the header "User Profile" displays a summary of the user model. This makes it clear that the personalisation is based on the values shown here. The user may click on the "Change your profile" link to update the values, and thereby change the user model.

If the user ever wonders how the system adapts the information delivered to him, he can click on the link that says "How was this page adapted to you?" at the bottom of the page, and a page like Figure 2 will be generated.

Parts that are highlighted with the lighter (yellow) colour were included. In this example, the first few lines were provided because of the user model values. They are the same as the first paragraph in Figure 1. Parts highlighted with the darker (green) colour were excluded. For this user, it was a much longer text, since his user model indicates that he requires a basic level of descriptions delivered to him. By placing the mouse cursor over the highlighted regions, the user may get a brief explanation, as a tooltip, explaining why the system included/excluded that part. Should the user be dissatisfied with the presentation, she can update her profile by clicking the "Change your profile" link.

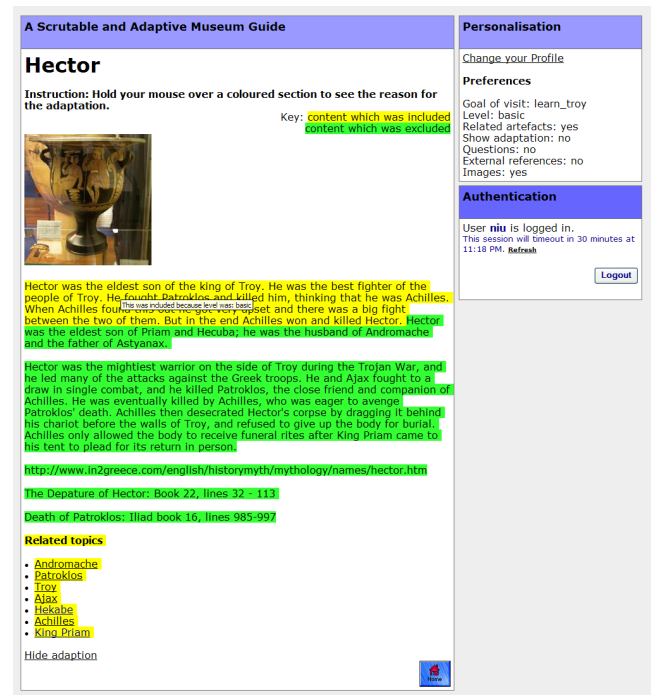


Figure 2: Scrutinising the Adaptation

3. SYSTEM ARCHITECTURE

This prototype system is based on a Web framework called Cellerator². This simple yet powerful authoring tool integrates well with Personislite, a lightweight version of Personis [5] which manages user models. The Web-based interface, adaptability, and controls for scrutability make it a suitable adaptive hypertext system for building up this prototyping.

²<http://www.it.usyd.edu.au/~bob/web.cgi?Cellerator>

User Modelling

The personalised information delivery is controlled by the user model. The current prototype uses a very simple user model with only seven components. Six of them are listed in Section 2, with one missing being a flag indicating if the user currently logged in is a first-time user. This component is used by the user modelling system to decide whether or not to initiate a new user model for that user. In earlier work [1], we have modelled users by exploring evidence of the duration of time the user has spent viewing or interacting with a particular digital object. In MyMuseum, this will map to the time spent at a museum exhibit. Based on this evidence, the user modelling part of MyMuseum can draw conclusions about user preferences and from this, tailor the delivery of information.

User modelling with Personislite allows adaptive systems to easily manage evidence for user models, and provide mechanisms to handle changes, noise, and other forms of reasoning under uncertainty. It utilises *resolvers* to conclude a value for each user model component based on this evidence. These resolvers are crafted by the system designers with scrutability in mind. For example, the simple point resolver uses the last most reliable piece of evidence to conclude a value. More complex resolvers can make knowledge intensive interpretations of all the evidence for a component (e.g. taking account of the timing of the evidence, the source, and its reliability).

At any time, users should be able to ask the system why an adaptation delivered by the user model was performed, and the system should respond with the interpretation of 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.

Exhibit Ontology

As discussed in Section 2, one of the core tasks of MyMuseum is to recommend a touring path through the museum. Essentially, this involves two tasks: finding artefacts that are similar to one the user has just enjoyed visiting or, in the case that the user does *not* enjoy the current artefact, recommending something that is different on the dimensions that affected the user's assessment with this artefact.

We are currently focusing on the reasons why users enjoyed an exhibit (or conversely, did not enjoy it) that are due to their interest in the aspects that the exhibit presents. This means that we focus on the semantics of the information presented in this exhibit. We are well aware that this is not the only factor in determining user satisfaction and interest. However, it is one crucial factor, especially in the case of museums which are visited by students as part of their formal education programmes.

The approach we are using to tackle this is based upon ontologies. Essentially, an ontology captures the meaning of concepts in terms of the relationships between them. So, for example, Paris is related to Troy because he was a prince of Troy. The core idea we are exploring is that if the user enjoyed an exhibit that was about Troy, they may be likely to also enjoy another exhibit that deals with Paris. To apply this approach, we need an ontology which captures relevant relationships. At one level, the simplest approach is to handcraft a suitable static ontology. This is the approach taken with MyMuseum. However, there are problems with this ap-

proach. First, it means that someone has to undertake the laborious task of creating the ontology. Even more serious is the problem that it is probably not possible to define a single static ontology that suites all purposes and all users. In particular, in the museum where we work, there are several major classes of users, including, for example, primary school students visiting as part of their school syllabus, final year school students of ancient history as well as university students in archeology courses. Taking the case of the primary school students, the important relationships between exhibits are defined by the knowledge of those students and the topics in their syllabus, as well as the particular areas their teacher has recently introduced.

To cope with this diversity of needs, we have taken an approach which automatically builds ontologies from existing documents, such as the primary school history syllabus and texts as well as the senior high school syllabus and resources on topics within that syllabus. Essentially, we aim to automatically build a dynamic ontology to drive the recommendation process.

From a theoretical view, this means generating a number of different ontologies from different sources (e.g. a dictionary, a textbook, a Website, etc.) for the exhibits. Depending on the visitor's interests and background indicated in the user model, one or more ontologies may be used to construct a suitable list of recommended artefacts to visit next. The visitor may scrutinise how the system decides the recommended artefacts during the tour, and she/he may update any invalid beliefs the user modelling system holds about her/him.

From the perspective of Personis, each relationship between two exhibits is modelled as a (exhibit model) component. Each component then has a list of evidence that states how the exhibits are related as well as shows the registered sources that reason about the relationship. For example, a component states that Paris is related to Troy. The list of evidence attached to this component is in the format of `<relationship>`, `<evidence source>`, `<scrutability support>` that may be:

- princeOf, High school history textbook, (parser1, URI1)
- returnTo, In2Greece³ online source, (parser2, <http://www.in2greece.com/paris.htm>)
- princeOf, In2Greece online source, (parser2, <http://www.in2greece.com/paris.htm>)
- destructorOf, Mythology Website, (MECUREO, URI2)
- sameContinentOf, Geography, (parser3, URI3)

The first word, separated by commas, describes the relationship between the terms (i.e. Paris and Troy). The second set of words specifies the source of this evidence, and the third set is a pair that provides a reference for scrutability control, namely what software has been used to extract this relationship and the original location of the source (e.g. a page number, an URL, a line number, etc.). The original location of the source is used to present explanations of the system reasoning.

As the relationship may be found multiple times in a single source, there could be multiple entries of evidence in the

³A Website about Greece that contains Greek mythology (<http://www.in2greece.com>, accessed 29 June 2005).

component. For example, Paris was a prince of Troy, and he caused the destruction of Troy. It is perceivable that some evidence may be more relevant to the visitor than the other. A high school student, for example, may be more concerned about the relationships of the terms that appear in the textbook than the ones in Mythology Website. Hence, when evaluating the relationship for a particular visitor, each piece of evidence may be assigned different weighting depending on the sources and the visitor's background and interests. In the case of finding a relationship between the prince Paris and the city of Troy, the evidence from Geography of Europe, which would probably mistake Paris as a city, could well be irrelevant. The visitor may scrutinise how the weighting is assigned for each source. If the user wishes, he/she can either change the weighting or update his/her user model.

4. EVALUATION

A small-scale, qualitative evaluation was carried out with seven participants, three women and four men. Two of them were experts in ancient history as they were in the midst of completing their honours thesis in this subject area; the rest of them had relatively limited knowledge of ancient history. None of them had used any sort of interactive museum guide previously, and all had a relatively low understanding of information technology. There were a total of 32 web pages used in this evaluation, with 28 of them being adaptive.

The participants were asked to use the system as if: 1) the system knew where they were looking, 2) the appropriate page(s) automatically appeared, and 3) they were in front of the artefact(s). Participants after they used the system then completed a questionnaire in three sections: awareness of scrutability, quality of information, and user interface.

The participants showed signs of not being used to the idea of scrutinising the adaptation. However, all users indicated that at some point they had examined their own profiles and made some changes in order to understand the difference between each of the settings. Regarding the quality of information presented, all users agreed that the adapted tour furthered their knowledge in the Troy era. Nevertheless, the response to the links and references to additional material suggested that these were mostly suitable for users with more background knowledge and interests. Hence, another level of content adaption seemed to be necessary to deliver better personalised information. In the user interface section of evaluation, the responses suggested some improvements were yet to be done. The main reasons were cluttered text on some pages and the obscure location of the "How was the page adapted to you?" link—at the bottom of the page.

Overall, the participants all found the system useful for a museum visit, although the "wow factor" [3] frequently posed a problem in terms of getting the participants to evaluate the intended usability issues.

5. CONCLUSIONS

A MyMuseum prototype has been built to explore ways to visualise how the system adapts to the user, as well as the reasons for the adaptation. From the evaluation carried out, it was clear that the concept of scrutability was not well perceived. While designing a system with scrutability in mind, we will also need to design it so that it is both

user-friendly and easy to perceive the concept.

One insight that was gained from the evaluation is the use of a dynamic ontology to deliver personalised recommendation of a visiting path. We will soon be implementing this onto MyMuseum using PersonisLite and expect to carry out another user study later this year.

6. ACKNOWLEDGMENTS

We would like to thank Smart Internet Technologies Co-operative Research Centre for funding part of this project.

7. REFERENCES

- [1] T. Apted, J. Kay, A. Lum, and J. Uther. Visualisation of learning ontologies. In *AIED '03: the proceedings of the 11th International Conference on Artificial Intelligence in Education*, July 2003.
- [2] G. Eckel. Immersive audio-augmented environments: the listen project. In *IV'01: The Fifth International Conference on Information Visualisation*, page 571, London, England, July 25 - 27 2001. IEEE Computer Society. ISBN: 0-7695-1195-3.
- [3] M. Fleck, M. Frid, T. Kindberg, M. Spasojevic, E. O'Brian-Strain, and R. Rajani. From informing to remembering: deploying a ubiquitous system in an interactive science museum. *IEEE Pervasive Computing*, 1(2):13-21, April 2002.
- [4] M. Hatala, L. Kalantari, R. Wakkary, and K. Newby. Ontology and rule based retrieval of sound objects in augmented audio reality system for museum visitors. In *SAC '04: Proceedings of the 2004 ACM symposium on Applied computing*, pages 1045-1050, New York, NY, USA, 2004. ACM Press.
- [5] J. Kay, B. Kummerfeld, and P. Lauder. Personis: A server for user models. In *AH'02: Proceedings of Adaptive Hypermedia and Adaptive Web-Based Systems*, pages 203-212. Springer-Verlag, London, UK, 2002.
- [6] P. Marti, A. Rizzo, L. Petroni, G. Tozzi, and M. Diligenti. Adapting the museum: a non-intrusive user modeling approach. In *UM'99: Proceedings of the Seventh International Conference on User Modeling*, 1999.
- [7] D. Petrelli, E. Not, M. Sarini, C. Strapparava, O. Stock, and M. Zancanaro. Hyperaudio: Location-awareness + adaptivity. In *CHI'99: Proceedings of Conference on Human Factors in Computing Systems*, 1999.
- [8] N. Proctor and C. Tellis. The state of the art in museum handhelds in 2003. In D. Bearman and J. Trant, editors, *Proceedings of Museum and the Web 2003*, Toronto, 2003.
- [9] E. Rich. User modeling via stereotypes. *Cognitive Science*, 3:329-354, 1979.
- [10] F. Sparacino. The museum wearable: real-time sensor-driven understanding of visitors interests for personalized visually-augmented museum experiences. In *Proceedings of Museum and the Web 2002*, Boston, April 17-20 2002.