

User-centred design of flexible hypermedia for a mobile guide: reflections on the hyperaudio experience

PETRELLI, Daniela and NOT, Elena

Available from Sheffield Hallam University Research Archive (SHURA) at:

<http://shura.shu.ac.uk/2906/>

This document is the author deposited version. You are advised to consult the publisher's version if you wish to cite from it.

Published version

PETRELLI, Daniela and NOT, Elena (2005). User-centred design of flexible hypermedia for a mobile guide: reflections on the hyperaudio experience. *User Modeling and User-Adapted Interaction*, 15 (3-4), p. 303.

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

This paper is a preprint of an article accepted for publication in UMUAI – User Modeling and User-Adapted Interaction The Journal of Personalization Research - special issue on "User Modeling in Ubiquitous Computing.", A. Krüger & A. Jameson eds. © copyright 2005 Springer
The original publication will be available at <http://www.springerlink.com/>

User-centred Design of Flexible Hypermedia for a Mobile Guide Reflections on the HyperAudio Experience

Daniela Petrelli^{a1} Elena Not^b

^aDepartment of Information Studies
University of Sheffield – UK
d.petrelli@shef.ac.uk

^bCognitive and Communication Technology Division
ITC-irst – Italy
not@itc.it

Abstract. A user-centred design approach involves end-users from the very beginning. Considering users at the early stages compels designers to think in terms of utility and usability and helps develop the system on what is actually needed. This paper discusses the case of HyperAudio, a context-sensitive adaptive and mobile guide to museums developed in the late 90s. User requirements were collected via a survey to understand visitors' profiles and visit styles in Natural Science museums. The knowledge acquired supported the specification of system requirements, helping defining user model, data structure and adaptive behaviour of the system. User requirements guided the design decisions on what could be implemented by using simple adaptable triggers and what instead needed more sophisticated adaptive techniques, a fundamental choice when all the computation must be done on a PDA. Graphical and interactive environments for developing and testing complex adaptive systems are discussed as a further step towards an iterative design that considers the user interaction a central point. The paper discusses how such an environment allows designers and developers to experiment with different system's behaviours and to widely test it under realistic conditions by simulation of the actual context evolving over time. The understanding gained in HyperAudio is then considered in the perspective of the developments that followed that first experience: our findings seem still valid despite the passed time.

Keywords: User Centred Design, Flexible Hypermedia, Mobile Guides, Content Adaptation, Development Support Environments.

1 Introduction

To guide the design of information systems towards actual user needs and expectation, human-computer interaction researchers have developed appropriate methodology and techniques. The user-centred system design (UCD) approach revolves around end-users. Potential users are involved from the very beginning and are regularly consulted for the evaluations of incremental prototypes (Preece et al., 2002). Though, a rigorous user-centred design does not start with a prototype, but with an extensive analysis of potential users, tasks, and environment (Hackos & Redish, 1998). Multiple techniques can be used and the analysis of the data collected

¹ The work discussed in this paper was carried out when the author was at ITC-irst in Italy.

should specify user requirements and system features. This starts an iterative process of user evaluation, redesign, and prototyping that ends when the system satisfies usability criteria (Nielsen, 1993; Harston, 1998).

UCD principles have been rarely applied throughout the whole design of adaptive systems. When adopted, user studies have affected the design of the user model, sometimes the interface layout (for example Bontcheva, 2001; Vassileva, 1996), but a pervasive user-centred design has hardly ever influenced the information organization or adaptation rules². Instead, a deep understanding of user, usage, and environment is instrumental in identifying what is the most appropriate content for each user class and can help in deciding where simple or complex adaptive mechanisms have to be applied. As a matter of fact adaptive systems can be implemented by using very simple techniques (e.g. triggers associated with users' actions) or highly sophisticated ones (e.g. deductive and inductive system reasoning) (Kobsa et al., 2001). Deciding about the needed complexity is a design decision and should be driven by the knowledge acquired on users and tasks during preliminary studies.

However, a good starting is not enough to assure the final adaptive systems would be user-compliant. UCD advocates an iterative process where incremental prototypes are developed and tested. Applying this principle in the context of adaptive systems requires the adoption of a modular architecture to support experimenting with different options. Indeed designing an adaptive system is not limited to working out a single solution; rather "the designer [of an adaptive system] specifies a number of solutions and matches those with the variety and the changeability of users and the environments" (Benyon, 1993). Conceiving different solutions implies for the designer a wide exploration of the range of alternatives in an iterative testing process. Moreover the more complicated the scenario the more difficult the exploration, given that adaptivity is then not limited to adjusting to users: factors such as where the action takes place, the device the person is using, and the communication infrastructure are all suitable subjects for adaptivity (Petrelli et al., 2001).

To assure that the adaptation is working as expected tests have to be done on real data. Indeed the effectiveness of an adaptive system can be judged only by assessing the actual format that is delivered to the user. In mobile and adaptive hypertext, for example, predicting how a page is composed at run time can be challenging: content and links included do not depend on the actual status of the user model only, but also on the current *interaction context* (where the user is, whether she is moving or not, what she is looking at,...). An extensive testing becomes mandatory to assure a smooth and coherent flow of information. Authoring data and extensive testing have then to be done in pairs. Although authoring support for adaptive hypermedia has always been considered important, only recently it has received the needed attention from both practical and theoretical perspectives (Calvi & Cristea, 2002; Cristea & Aroyo, 2002; Weber et al., 2001; De Bra et al., 2003). Still, data creation and rules testing are kept apart, possibly because content creation is considered the task of domain experts while rule testing is developers' responsibility. When the scenario of the interaction is not limited to screen, keyboard, and mouse, as in the case of mobile guides, an environment for testing how each context component contributes to the final adaptation is a valid support for system development. Designers of adaptive systems would benefit from a tool that supports fast prototyping and testing of new promising ideas. The same environment should then be used to produce the annotated data and for testing its adapted form as delivered to the user.

As discussed above, applying UCD to the design of adaptive systems is particularly challenging because the behaviour of the final system is intended to dynamically adjust according to multiple parameters, i.e. user preferences, knowledge and behaviour, and interaction context.

When, in the mid 90s, we first started working on one of the first prototypes of adaptive and mobile museum guide (called HyperAudio, Not et al., 1997a), not much experience was available in the Adaptive Hypermedia community on how to export principles of adaptivity to mobile applications, nor much skill was available on the application of UCD to adaptive systems. In the initial critical phase of the project we faced problems like envisaging credible scenarios of use, identifying parameters for adaptivity, designing content and adaptation rules in a suitable way. The initial aim we had in mind was to offer the visitor personalized information centred on his/her current standing position. The envisaged interface was a web-based layout with an active involvement of browsing users. What the final development of HyperAudio offered instead was an experience of freely moving in an information space and automatically receiving tailored information. We

² Exceptions are web-based recommendation systems that make use of massive logs of user profile/behavior to select the most appropriate information and in general to implement adaptivity (Kobsa et al., 2001).

started with the idea of an adaptive hypermedia displayed on a PDA for browsing and ended with an intelligent system that was responsive of social and relational conditions, of visiting pace, of visitor's interests. It was intended to be a guide; it ended being a companion.

This deep change in how the adaptation should be manifested was due to an extensive survey of museum visitors coupled with an explorative design, as explained in the rest of this paper. The deep analysis was not limited to descriptive statistics (e.g. the percentage of visitors who arrive at the museum already informed) but also included correlating different data (e.g. those more likely to be families) and ultimately designing solutions (e.g. consider families as a separate user class). Results supported the decision to go for a simpler and lighter architecture but a more sophisticated data structure than originally conceived.

The experience we gained in the small scale HyperAudio project contributed ideas to HIPS, Hypernavigation In the Physical Space (Benelli et al. 1999), a broader European project funded in the i3 (Intelligent Interactive Interfaces) framework, where we further explored the UCD approach by creating a workbench for fast prototyping and off-line testing. The use of such development environment closed the cycle of UCD applied to adaptive systems: we could test different solutions by simply "plug-and-play" different modules (e.g. different user models, different adaptation rules), and we could verify the system was behaving (i.e. adapting) as expected by performing a set of off-line tests.

Since HyperAudio initial implementation, many other systems have been developed according to the principles of seamless and personalized interaction with the surrounding space (see section 7 for some references), however, as an insight in the evolution of design, we consider our experience still valuable and unique after all these many years. This paper reports on HyperAudio experience of applying the UCD approach when developing a handheld museum guide that adapts its behaviour to users, their position in the space, and their interaction with both the guide and the environment. The architecture of the HyperAudio system and its sophisticated adaptive mechanisms are discussed in section 2. The user study conducted to understand Science Museums and their visitors is described in section 3; the redesign of the first ideas of user model, data structure, and adaptation rules follow in section 4 together with some scenarios of use. A discussion of the importance of an interactive environment for fast prototyping and component testing for design purposes follows in section 5, while section 6 discusses the use of the same environment for data editing. Finally, section 7 presents related work in immersive and adaptive mobile guides.

2 HyperAudio: Location Awareness and Adaptivity

2.1 The History

The late 90s saw a substantial increase in the work done in adaptive hypermedia towards the most diverse domains (Brusilovsky, 2001). That was also the time when the idea of adapting an existing hypertext to the interacting user by means of a user model came into contact with research into natural language text generation. Researchers in natural language processing were developing dynamic hypertext, where pages are generated on the fly on the basis of some domain knowledge representation, according to a user model (Oberlander et al., 1998; Milosavljevic et al., 1996). The First Flexible Hypertext Workshop (Milosavljevic et al., 1997) was a forum for discussing and comparing the two approaches and other hybrid solutions.

At the same time, the human-computer interaction community was exploring the new world of mobile devices (Johnsons, 1998). The ideas of augmented reality and ubiquitous computing of the early 90s (CACM, 1993) were maturing into exciting experimental systems able to locate the user's position via sensors, and to react accordingly, e.g. by switching on/off electronic devices or transferring data to support the user's task (Bederson, 1995; Abowd et al., 1997).

Our project started in 1997 with the aim of fusing these hot issues in the areas of Adaptive Hypermedia, Natural Language Generation and Human Computer Interaction. The challenge was to create a *smart* location-aware system for delivering personalized hypermedia to an itinerant user. Museums were chosen as a promising application test-bed because visitors move in the physical space looking for interesting exhibits and wishing to acquire information, deepening their knowledge and interests. The museum was then thought as a sort of augmented environment, sensitive to visitors' movements, where an information hyperspace can be associated to each exhibit; the visitors explore that hyperspace during the physical visit (Not et al., 1997a,

1997b). The envisaged system would automatically play an audio comment as soon as the visitor approached an exhibit. Since the main communication channel was intended to be audio and the information was presented with a hypertextual paradigm we chose *HyperAudio* as the project name (Not et al., 1998). However prominent, audio was not intended to be the sole presentation medium: a dynamically created hypermedia page would display images, text, and links potentially interesting to the visitor. The presentation (audio message and hypermedia page) would be adapted to the interacting user but also to the broad interaction context, including the physical space, the visit so far, the history of interaction, and the narration.

2.2 The Challenge

In the HyperAudio project we interpreted the term “adaptation” in its broadest sense. The system had to adapt its behaviour to serve at best the visitor’s goal of enjoying the museum, finding it rewarding and useful. Thus the system had to adapt the presentation content and navigation hyperlinks to each visitor, but also had to take into account the prominence of the physical space, the objects of interest, and the visitor’s position with respect to them. The guide had to select content about the object in sight or apply strategies to attract the visitor’s attention towards other objects. Moreover what the visitor has already seen (in the physical space) or heard (from the hypermedia space) had to be considered: a properly designed adaptive guide would not propose the same information twice to a visitor coming back to an already seen object.

However, the user model, the space model, and the visit history were not considered sufficient for assuring a smooth interaction with HyperAudio. The sequence of messages delivered to the user had to be a single smooth narration, thus the composition of the presentation had to consider rules for effective content structuring and linguistic realization according to the current discourse context. For example, during a presentation, a deictic reference to an object in the physical space, like “*this is* the fossil of an ancient crocodile”, is valuable to reinforce coherence between vision and text. On the other hand, other appropriate lexico-grammatical patterns may be used to manifest certain kinds of semantic relations between text units which reinforce *texture*, i.e. the property of a text of being perceived as coherent (Halliday & Hasan, 1985). This happens for example when appropriate anaphoric referring expressions are used, like the pronoun “it” in “it was found under a thick rock stratum.”, or when markers are used to make explicit the rhetorical relations between content units, like “conversely” in the following example “Reptile skin is covered by keratin or horn scales. Their position and thickness prevent desiccation. *Conversely*, amphibians have naked skin that lacks protective devices.”.

The overall HyperAudio challenge was therefore to select the most appropriate content and links with respect to the current visitor’s interests, the environment, and the interaction so far, and to polish the final presentation by adjusting the narration. The following section discusses the adopted solution.

2.3 The Hardware and Software Architecture

In the HyperAudio scenario, the visitor is provided with a palmtop (an Apple MessagePad) equipped with headphones on which an infrared receiver is mounted (Figure 1). Visitors are asked to position the infrared receiver under their chin, in order to ensure that only signals coming from the front are detected. Each meaningful physical location (e.g., exhibit, door, or passage) has a small power-autonomous infrared emitter that continuously sends a uniquely identifiable code. Thus the physical space is partitioned into sensitive zones allowing the system to identify the visitor’s position and orientation (the *Space Model* in Figure 2).



Figure 1. HyperAudio hardware.

When the user approaches an exhibit, the corresponding infrared signal is detected (*implicit input*), the system is triggered and a description (*presentation*) of the object in sight is dynamically composed. The presentation has an audio message and an image relevant to the object described, plus a set of suggested links. The pointing of the pen on a displayed link (*explicit input*) activates the system as well, as outlined in Figure 2.

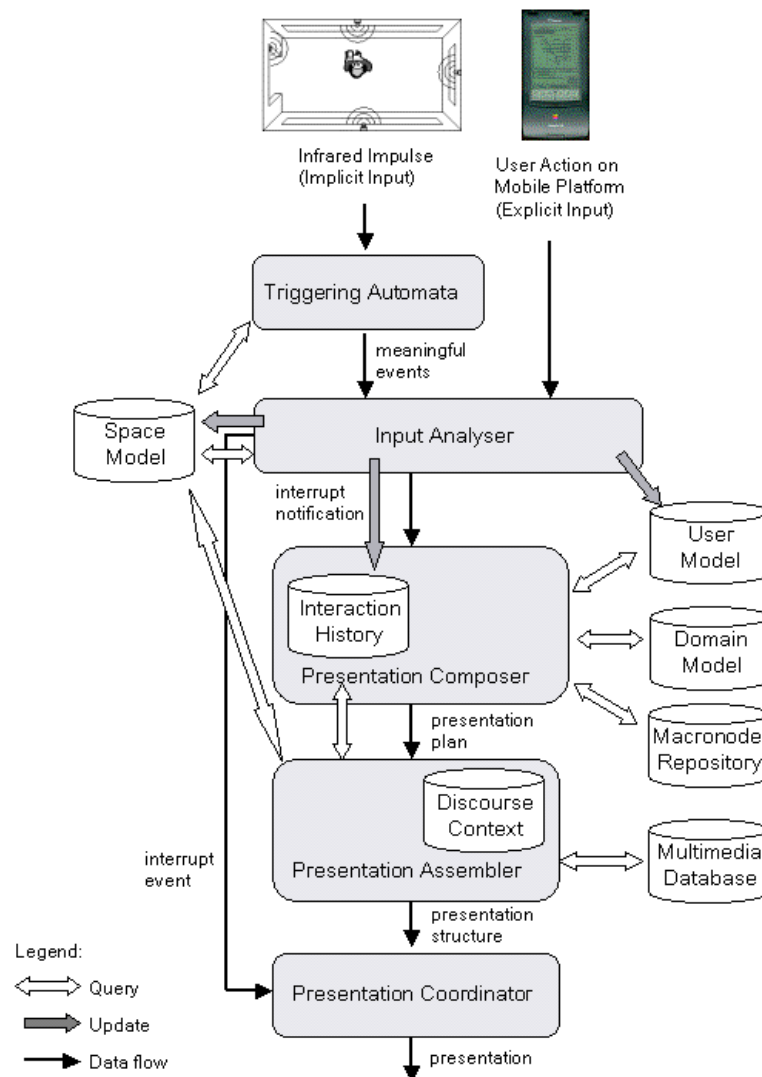


Fig. 2 The HyperAudio software architecture.

While the selection of a link on the interface is clearly a request, the implicit interaction generated by a movement has to be validated. When an infrared signal is detected, the *Triggering Automata* queries the *Space Model* for the user's current and previous positions. By comparing the two it determines whether the user entered/exited a sensitive area and how much time she spent in each cell of the augmented environment; quick changes of positions (i.e. different signals received in a fast sequence) are discarded as noise, others are passed on to the *Input Analyser* as *meaningful events*.

The *Input Analyser* decides the most appropriate reaction to implicit and explicit inputs. For example it sends an *interrupt event* to the *Presentation Coordinator* if a different cell is entered (i.e. the visitor significantly changed her position) and asks the *Presentation Composer* for a *new presentation* appropriate to the new position. It also updates the *User Model* following the visitor's actions; for example, stopping a presentation shows negative evidence and the interest model is updated accordingly.

The *Presentation Composer* is responsible for planning the overall presentation that integrates (where appropriate) object descriptions, images, links for follow-up information requests, and oriented maps. To create a *presentation plan* the composer traverses an annotated multimedia network stored in the *Macronode Repository* and uses the knowledge in the *Domain Model*, the *User Model*, and the *Interaction History* to

decide which nodes will be included in the audio presentation, which will become links and which will simply be ignored (see sections 2.3.1 and 2.3.2). The sequence of presentation plans is cumulated in the *Interaction History*, which keeps track of what has been presented to the user so far.

The *Presentation Assembler* takes the presentation plan and actually assembles the final message. Here is where the linguistic arrangement takes place with respect to the current *discourse context*, as explained in 2.3.2. Finally the assembler substitutes symbolic names with the appropriate multimedia data (audio files, images and maps) and asks the *Presentation Coordinator* to physically deliver the presentation to the user.

2.3.1 The Annotated Data Structure: Macronode Formalism

Adaptive hypermedia are based on the idea that pages and links are appropriately annotated so that personalization can be computed at run time. The amount and the type of annotation depend on the system (Brusilovsky, 1996). For example an HTML page can be annotated with structured comments that indicate when a piece of information (text or link) has to be included (De Bra & Calvi, 1998). However dynamic hypermedia do not keep any underlying network, but generate each page on the fly from some knowledge representation (as in ILEX (Oberlander et al., 1998)). The solution adopted in HyperAudio is hybrid: there is a richly annotated network of information units from which presentations are built, but nodes do not correspond to pages but rather to fragments of a page (Not & Zancanaro, 1998). Strategies borrowed from the field of natural language generation are used to select and structure information units and properly assemble multimedia pages (where audio plays a major role), adjusting the linguistic realization of the message to guarantee coherence and cohesion of the final presentation.

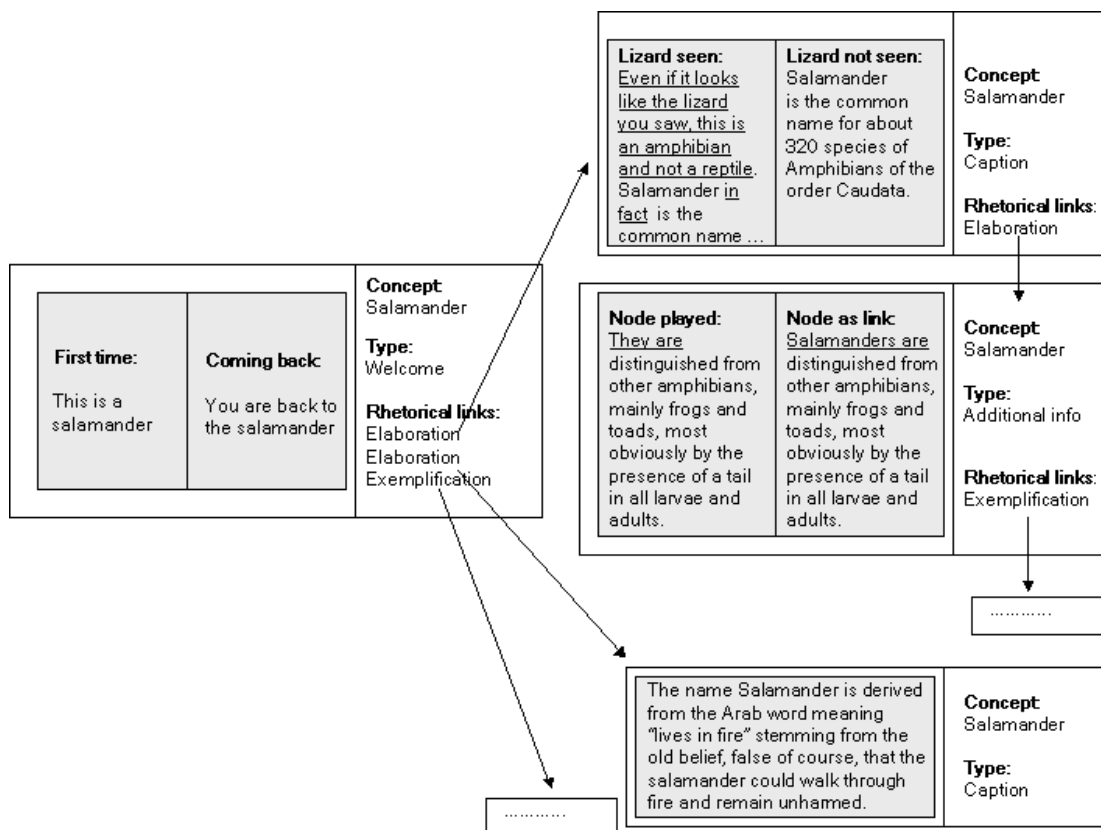


Figure 3: A network fragment of (simplified) Macronodes: in the content part (left) linguistic variations are underlined.

Content selection is enabled by the fact that each information unit, which we call a *macronode*, contains a shallow semantic annotation that describes its main topic (i.e., what the node is about) and its function in the

narration (i.e., introduction, core information, or additional details). Macronodes in the repository are related to each other by rhetorical relations (Mann & Thompson, 1988) that help describe the semantic relations between the various information units and how they could be textually integrated coherently. A macronode is internally organized to allow for some linguistic variation. Figure 3 shows a sample fragment of macronode network. The linguistic adjustments are actually computed at run time by the Presentation Assembler which selects from a conditional graph (see Figure 4) the most effective realization according to constraints over the space model, the discourse context and the interaction history.

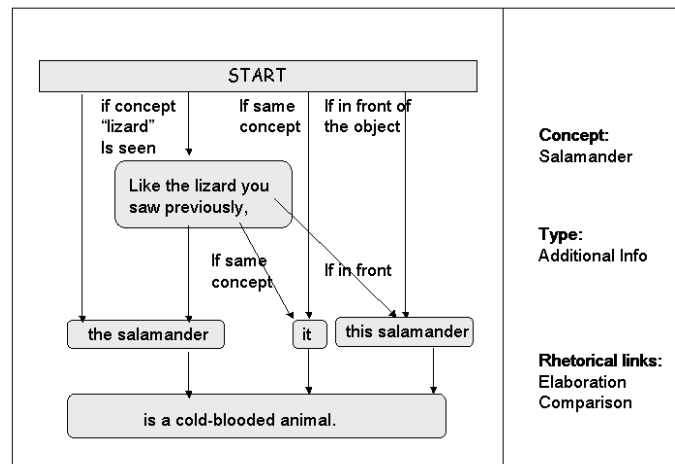


Figure 4. The macronode internal structure.

The content of the macronode shown in Figure 4, for example, could result in the following alternative sentences:

- “like the lizard you saw previously, the salamander is a cold-blooded animal”
- “like the lizard you saw previously, it is a cold-blooded animal”
- “like the lizard you saw previously, this salamander is a cold-blooded animal”
- “the salamander is a cold-blooded animal”
- “it is a cold-blooded animal”
- “this salamander is a cold-blooded animal”

In the original implementation of macronodes, all linguistic adjustments of the macronode’s surface form were realized through conditional text, manually specified by the content author with the aid of a macronode editor (Petrelli et al., 2000). The text was then to be read and recorded by a human actor. Further research has investigated the integration of this manual approach with the automatic generation of sentences or portions of sentences (e.g. the insertion of pronouns or deictic references, or the reference to previously seen objects), to relieve the author’s burden when a speech synthesizer is available (Not & Zancanaro, 2001).

2.3.2. The Adaptation Techniques: Input Analyser, Presentation Composer, and Presentation Assembler

As described above, HyperAudio has three points where adaptivity is realized. Different sets of rules are used by the different modules for deciding (i) if a presentation has to be composed, (ii) eventually composing it, and lastly (iii) tuning its final linguistic form.

The first set is used by the Input Analyser and includes rules such as “if the visitor is leaving an object, then interrupt the running presentation” or “if the visitor approaches a new object but the current presentation is general, then let it finish”.

Rules applied by the Presentation Composer decide about content and links selection, as well as the length and the inclusion of new concepts. Strategies are encoded to avoid presenting already known information, to choose the kind of information for which the user’s interest is high, to present new information when the user goes back to a previous topic. In addition, rules checking the rhetorical links between macronodes control the

length of an elaboration chain or the inclusion of background information to clarify a topic.

The Presentation Assembler takes care of tuning the linguistic form of the presentation considering the current status of the Space Model, the Discourse Context, and the Interaction History; it applies rules such as “if the user is in front of the object, then select the text containing a deictic reference (e.g., ‘this is’)” or “if a concept has been already introduced (e.g., the object has been seen), then include an explicit reminder (e.g., ‘you saw previously’)”.

2.4 User Interface

The design of the interface was based on two basic constraints: (i) the MessagePad screen has low definition; and (ii) visitor’s attention is devoted to the exhibition. As a consequence, the audio channel mediates the descriptions of the exhibits whereas the graphical interface is reduced to the minimum. Figure 5 shows a typical screenshot: a central picture provides the context of the current description, and links to concepts related to the object in sight are displayed as buttons. Those above the picture lead to other related concepts, those below the picture lead to elaborations of the same concept. By clicking on the buttons the user can explore concepts related to objects located elsewhere in the exhibition. A map, displayed by clicking a further button, shows the position of the object currently described, whether in the room or elsewhere. Finally a “back” enables repeated listening to previously played presentations.



Figure 5. The interface layout as displayed (a) in the reptiles room, and (b) in front of the lizard.

3. User Requirements Elicitation

As mentioned in section 1, UCD typically starts with an extensive analysis of potential users, tasks and environment that feeds the design of the first prototype. However given the novelty of the topic we decided to test the actual feasibility of the system at the same time as the user study was going on. A first functional prototype was implemented in Spring 1997 as a proof of concept. A user study was then set up to elicit user requirements and obtain ideas for the design of the user model and adaptation rules. The main purpose was to identify user characteristics that would compose the profile to be implemented as stereotypes (Rich, 1989). That seemed the best choice considering the constraints of using a PDA: adaptation had to be simple and light, quick and effective from the very beginning.

Visitors' behaviour has been studied for long and a whole museum literature is devoted to this topic. However this extensive knowledge was of limited help in defining stereotypes and adaptivity, since the focus is generally on how exhibition layout affects moves and how to make it effective. Since data on how personal traits relate to behaviour were not available, a user study was set up. We hypothesized that visitors' behaviour could be predicted using "classical" dimensions such as age, profession, education, specific knowledge or background. This categorization would allow setting features in the user profile, such as language style (expert vs. naive) or preferred interaction modalities (led by the system vs. led by the user). The study was not intended to be a survey of museum visitors; personal data were of interest if (and only if) they correlate with a predictable behaviour. The objective was to discover, for example, whether aged people have a negative attitude toward technology and would prefer to be guided; this would militate for a non-interactive setting of HyperAudio; conversely, a positive attitude expected from younger people would motivate a highly interactive mode.

3.1. The Case Study

To find out if relations between personal facts (e.g. age, specific interests) and the way museums are visited do exist, a survey was conducted. A questionnaire was organized around five topics:

- **A personal data profile section** asked for demographic information: Age, sex, education, job, and residence were selected as having potential for discriminating attitudes towards museums.
- **A museum habit section** complemented the personal profile. It collected data on how often the respondent visit museums and what are their preferred ways of visiting, e.g. alone, with a partner, with the family, or within a guided tour.
- **A context of the current visit section** focused on the just finished visit and asked if it was the visitor's first time in the museum, with whom they came, and the general motivation for the visit.
- **A course of the current visit section** collected opinions on the use of guides (from books to human guides) as well as the duration and the purpose of the visit.
- **A styles of visit section** collected general attitudes and opinions on different ways of visiting museums.

The final version of the questionnaire was composed of 26 questions, tested by a pilot study. It required around 10 minutes to fill in and was introduced by a page describing the purpose of the study. The survey was conducted from October to December 1997 in three museums focusing on topics related to natural science. As they exited, visitors were asked to take part in the study by museum staff and a total of 250 answers were collected.

3.2. Discovering Visitors' Attitudes

Empirical results revealed relevant and unexpected facts that required the designers to rethink their initial assumptions. The main findings are summarized in this section (for details see Petrelli et al., 1999).

The most unexpected and disappointing outcome was that personal data, like age, profession, education etc., did not account for respondents visit attitudes. Older people did not show preferences different to those of younger ones; education was high for almost all museum visitors (91%); professional interest had no impact. Thus, personal data would not predict visitor's behaviour and would not help in the adaptation process. As a consequence, asking for personal details in the initial form is of no use. Fortunately, other attributes were discovered which accounted for visit and interaction variability.

The social dimension emerged as an unpredicted important factor. Only 8% of visitors like to visit the museum alone; 24% prefer friends and partner; 20% choose organized tours; and a big 42% go with the family. Visiting a museum is mainly a 'social event' and being in a group changes the visiting pattern. Indeed, our data confirm Falk and Dierking (1992): people tend to behave differently when visiting museums with friends or family. When visiting museums with friends, adults are mainly concerned with the nature and content of the exhibits. Even if discussion is stated as a very important point, their attention is more focused on what they see than on their own social group. Conversely, adults with their family typically focus on their children, on making the exhibition understandable and the visit enjoyable. Family visits are led by children, and family learning (i.e., when adults and children learn together) derives from discussions (Borun & Dritsas, 1997). Our data showed also that families are more likely to arrive at the museum already informed than adult

groups. This indicated that classes of users had to be considered, with their different needs, expectations, and behaviours.

Another surprise was the number of non-first-time visitors, accounting for 68% of the sample. Being a frequent visitor was related to the type of visit and to the time spent visiting. Returning visitors came to see specific objects and stayed in the museum longer than those who came for the first time and wanted to see the museum in general. From this perspective the same behaviour may have a different meaning, e.g. skipping an object may indicate lack of interest, but this may not be the case for frequent visitors who have seen the object before. Thus a long-term user model (e.g. some kind of profile stored between visits) would be useful in such a context.

Visitors have a positive attitude towards guidance and use it if available (58% of our sample used a guide during their actual visit), regardless of personal attributes (e.g. age or knowledge). What accounts for the use of a guide seems to be familiarity with museums: the more visitors are used to going to museums, the more they use a guide. In addition, those who came to see specific objects used a guide, while people who came to visit the museum in general did not. These results are counterintuitive; we expected that familiarity with museums would reflect an autonomous self-sufficient style.

To reinforce the previous finding, only 7% reported liking using technological devices as museum guides. Most people liked visits guided by a member of the museum staff (53%), while 21% of the sample preferred catalogues or books, and 19% preferred to visit the museum without any support. These data led to several important considerations. First of all visit aids are highly appreciated. Secondly, the preferred solution is still human experts. This may be due to social factors and to the possibility of interacting with a source of knowledge, but it also suggests that listening to a human guide is still the easiest way to get information. Finally the general dislike of technology suggests that some visitors may never explicitly interact with the system. This completely passive behaviour of some users has strong implications, therefore the possibility for the system to provide a completely automatic visit was considered.

3.3. User Requirements and System Design

The survey study provided a deeper understanding of which are the important aspects of visiting a natural science museum. From this knowledge a set of user requirements were extracted:

- **Museum visit is a social activity:** groups have to be accommodated as well as single visitors.
- **Families (and schools) are important targets** and must be considered as distinct classes of users.
- **Families behave differently from adult groups:** families arrive with some background knowledge, the visit is driven by the children and learning comes from adult-children discussion.
- **Frequent visitors are important targets** and must be considered as a class.
- **Frequent visitors behave differently from first time visitors:** they see fewer objects and stay in the museum longer; this behaviour has to be accommodated.
- **First time visitors want a general overview:** they are not interested in details and have to be engaged if they are to return.
- **Attention is devoted to the exhibit or to the group** and not to the computer: the interaction has to be reduced to the minimum.
- **Guidance is welcome.**
- **Technology is disliked.**

The list was very different from the one expected, one where personal details would account for visiting attitudes; it became a tool for driving the interaction design and for generating new ideas. The anticipated interaction was also reconsidered. Before the study, the envisaged interaction was browser-based with text, image and links dynamically selected and composed; the audio message would direct the user's attention towards the PDA. Discovering that guided tours are well accepted and, more important, that interacting with technology is not a favoured activity changed our view. In this context³, a system that autonomously decides

³ In other scenarios this principle may not hold and control over the adaptive mechanism may be appropriate; however each solution has its own advantages and has to be considered in respect to each application (Jameson & Schwarzkopf, 2002). For example, when a proactive adaptive system is used to support activities in a daily working environment,

what to do (i.e., a self-adaptive system, Dietrich et al., 1993) was expected to have a greater appeal than one that asks for the user's assistance (a user-controlled self-adaptive system, Dietrich et al., 1993). This design decision seems also supported by Cheverst et al. (2002) findings that during the evaluation of GUIDE, the vast majority of users wanted to invest as little effort as possible in navigating for the retrieval of information. The HyperAudio final prototype supported also a proactive modality that automatically provides information, thus allowing for no interaction at all. Although a formal user evaluation never took place⁴, we observed many people using HyperAudio in a small museum simulation installed in ITC-irst: all were impressed by the reaction of the environment to their movements and virtually nobody took any notice of the device they were carrying. We had implemented the idea of *information appliances*, small devices dedicated to a single task (Norman, 1998): the action of visiting is kept as natural as possible and interaction with the computer disappears.

3.4. On Results Generalization

In the follow-up experience we did together with the European partners of the HIPS project, a wider study on user requirements was conducted (Broadbent et al., 1998). Besides questionnaires distributed to visitors, focus groups with stakeholders (e.g. museum curators, art experts, custodians) were held to more precisely depict the needs of both visitors and managers. The study was carried out in four different places, three different countries (Norway, Germany, and Italy), and focussed on art museums (modern art in Norway and Germany, historical palaces in Italy⁵). The goal of the questionnaire in this study was not precisely the same as in HyperAudio: HIPS questionnaire was focussed on which art features visitors appreciate⁶ more than on discovering actual behaviour and attitudes. However, despite the differences in the two questionnaires some degree of comparison is possible⁷.

The first sensible difference is in the type of visitors (Table 1) with a strong dominance of family and group in the science museums versus partner and friends in art museums.

	Alone	Partner	Friends	Family	Group
HyperAudio	8%	14%	10%	42%	20%
HIPS	12%	39%	25%	19%	6%

Table 1. Preferred visit companion.

A second difference is in the preferred guide (Table 2 summarises the data). While human or audio guide account for the majority of preferred supports in science museums, participants definitely preferred a more autonomous visit in art museums. Interestingly in both studies technological supports (the audio guide) were equally disliked.

	Maps	Guidebook	Leaflets	Human	Audio	Desk	Friends	None
HyperAudio		20%		52%	7%		6%	12%
HIPS	25%	21%	5%	13%	4%	3%	10%	21%

Table 2. Preferred visit support (not all the options were included in both questionnaires; multiple choice

even the possibility for the user of scrutinising/modifying the inner user model and system inference rules might be important as discussed by Cheverst et al. in this issue (Cheverst et al., 2005).

⁴ By the time the prototype was ready the MessagePad was no longer being produced or supported by Apple, thus the planned porting in the museum setting never took place.

⁵ The museums were (websites assessed 30.9.2004):

- Edvard Munch Museum in Norway – Munch's paintings and drawing <http://www.munch.museum.no/>
- Kunst Museum Bonn in Germany – modern and contemporary paintings <http://www.bonn.de/kunstmuseum/>
- Castello del Buonconsiglio in Trento (Italy) – a noble XIII-XVI palace <http://www.buonconsiglio.it/>
- Palazzo Pubblico in Siena (Italy) – the medieval townhall <http://www.comune.siena.it/museocivico/>

⁶ This was done to feedback with realistic data the design of the user interest model.

⁷ The original data is no more available to perform the same analysis done for HyperAudio and see if behavior and attitudes in art museums differ from those in science museums.

was allowed for the HIPS one and each % is calculated respect to the total number of answers).

Both tables above clearly show how the two contexts (science and art) are different and underline how assumptions based on somebody else's result can be risky. For example, caring particularly for family or group visitors seems not justified in art context. The need for a direct investigation on the reality is reinforced by some findings in the HIPS study: the questionnaire deeply considered the many facets of interest in art (e.g. technique, composition, theme, artist, social or political context, history, etc.) and showed a contrasting polarization of interests between the two clusters, namely historical vs. modern.

Although some results can be generalized and imported, e.g. the suggestion of including maps or guided tours in museum mobile guides (Broadbend et al., 1998; Broadbend & Marti, 1998), the information that can influence adaptation needs to be collected by the designing team and targeted toward the open questions that need direction. Undeniably the user requirements elicitation done for HyperAudio provided fundamental rationale for designing the adaptive behaviour, and for defining the appropriate internal data structures and adaptive rules.

4. System Redesign Based on User Requirements

Empirical evidence is used in UCD to direct redesign and adjustment. In HyperAudio this meant reconsidering the functionalities and adaptive behaviour the system was to support on the bases of the requirements collected in the previous phase. By analysing the requirements list we recognized how much of the flexible behaviour of the system could be implemented by simple adaptation techniques, like explicit triggers, instead of more complex reasoning. The user model, the data structure, and the matching rules, were revised from this perspective, as discussed below.

4.1. From a User Profile to a "Visit" Model

The strongest effect of the user study was on the user model. The original design of the User Model was based on a thorough study of the existing literature on how visitors typically behave in museums (e.g. Falk & Dierking, 1992), conversations with museum curators, and studies on how exhibition layout affects visitors' behaviour (Lozowski & Jochums, 1995). We intended to keep a detailed user profile to be collected via an initial detailed questionnaire. The questionnaire data was intended to be integrated with predicted attractive and holding power of each exhibit⁸ and used to initialize a model of the user's background knowledge, interest, and interaction preferences (Sarini & Strapparava, 1998).

As soon as the analysis of the user studies became available it was clear that some of our initial hypothesis about interest and knowledge modelling had to be revised. The idea of stereotypes based on personal features was abandoned while others were considered. Actually there has been a shift from a user model to a *visit model*. In relation to this, four features emerged as important and were included in the initial questionnaire:

- **Family, school or adult(s):** the three groups are different in interests, previous knowledge and ultimate goals. By knowing which group a visitor belongs to, the system can select different content for the presentation (e.g. classification vs. curiosity), can adopt a specific presentation style (e.g., narration vs. question-answering), and can automatically set an interaction mode (e.g., interactive for families, automatic for schools and adults).
- **First-time visit:** first-time and frequent visitors are differentiated. This affects content selection as well as the length of the presentation. For first-time visits the preferred content is introductory, actually an overview, while for subsequent visits a deeper content is preferred. The fact that frequent visitors spend more time and see fewer objects motivates the decision to use this information for setting the interest model to high so that longer presentations are composed from the very beginning.
- **Anticipated visit duration:** the more time is available, the broader the visit can be. This affects system verbosity in terms of numbers of objects proposed or depth of descriptions.
- **Interaction preferences:** proactive behaviour is the default mode; however it is considered

⁸ Attracting, and holding power are the probability that the visitor stops and observes the exhibit, and the average time spent by visitors in front of it (Lozowski & Jochums, 1995).

important to allow visitors to change from fully automatic (i.e. the system plays the message automatically as soon as the visitor reaches an active area) to interactive (i.e. the system announces that new information is available with a “beep” but it is played only when the user explicitly clicks) since this is a preference that cannot be easily inferred.

The neutral nature of those questions would allow museum staff to fulfil it on behalf of the user when the guide is handed out, thus providing personalized information also to passive visitors, people who would never explicitly interact.

The dynamic part of the user model was also revised. Initially conceived as a complex weighted activation network over domain concepts (Sarini & Strapparava, 1998), user interest was finally implemented as an array of boolean values, each item associated to a concept: an item is set to true for returning visitors or when visitors stay in front of the corresponding object for more than two seconds after a presentation has finished; it is reset when the presentation is stopped. The user knowledge model simply ticks already heard macronodes: when the Presentation Composer traverses the network for collecting macronodes for the new presentation those already heard are discarded. Finally a boolean value to regulate system verbosity was introduced; it would be on for long visits or returning visitors. This very simple implementation of the dynamic user model had the advantage of making computation efficient even with limited computational power as when using only a PDA. As a consequence the HyperAudio reaction to user moves was very fast and the natural pace of visit was not affected by the system.

4.2. Revisiting the Data Formalism

The macronode formalism discussed in 2.3.1 was refined on the basis of the results of the survey. In particular the fact that visitors belong to different groups with different goals, e.g. families vs. adult groups, suggested a richer information space and a finer description of the node content. The *perspective* field was added to the macronode structure in order to better describe how the main *concept* of the macronode was elaborated in the content unit (classification, curiosity, characteristic...). Adaptation rules would then prefer different macronodes with different perspectives on the same concept depending on the selected user class, as shown in the scenarios in section 4.4.

A broader range of text types was introduced as a further data refinement. The purpose was to better support frequent visitors in the in-depth exploration of a limited number of objects. Thus a distinction was made between linked information that must be played immediately (e.g. for frequent or interested visitors) and elaborations that can be added to the message or included as links (for an example see section 4.4).

A further alteration to the original data structure was the *presentation style* to distinguish different forms, such as narrative, question-answer or dialogues. As before, a different style can be associated with a user class preferring narration for adults, question-answering for families and dialogue between characters for pupils.

4.3. The Adaptive Rules

As discussed in 2.3.2, HyperAudio adaptivity is realized by different sets of rules directed to different objectives. All rules were revised as a result of considering the found requirements. Discovering that visitors might never interact suggested reinforcing system reactivity to physical actions; for example “if the presentation has finished and the visitor does not move for a further two seconds, then prepare a new presentation”.

Rules applied by the Presentation Composer were the subject of more revision. For example, the requirement to engage first time visitors suggested the rule “if a first-time visitor has a long time available, then propose visiting a new exhibit related to the current one”. It is worth observing that these rules are designed on the bases of few context elements (mainly the questionnaire and current interaction) but provide a wide range of flexibility.

New composition rules were also derived from the revision of the data structure described above. For example “if a frequent visitor, then choose the longest chain of macronodes available for the topic”, or “if family, then prefer a ‘curiosity’ perspective for the concept in focus” (similarly prefer ‘characteristic’ for adults and ‘classification’ for pupils). It should be noted that the association between user class and a specific

perspective is based on the inferred user intention, having fun for the family, learning basics for pupils, and acquiring generic knowledge for adults. However this association is quite arbitrary: different perspectives and associations would have been equally valid.

4.4. The Resulting System Behaviour

The following scenarios exemplify particular cases of visits and show how the macronode network shown in Figure 3 can be instantiated for different presentations using the rules that emerged from system redesign.

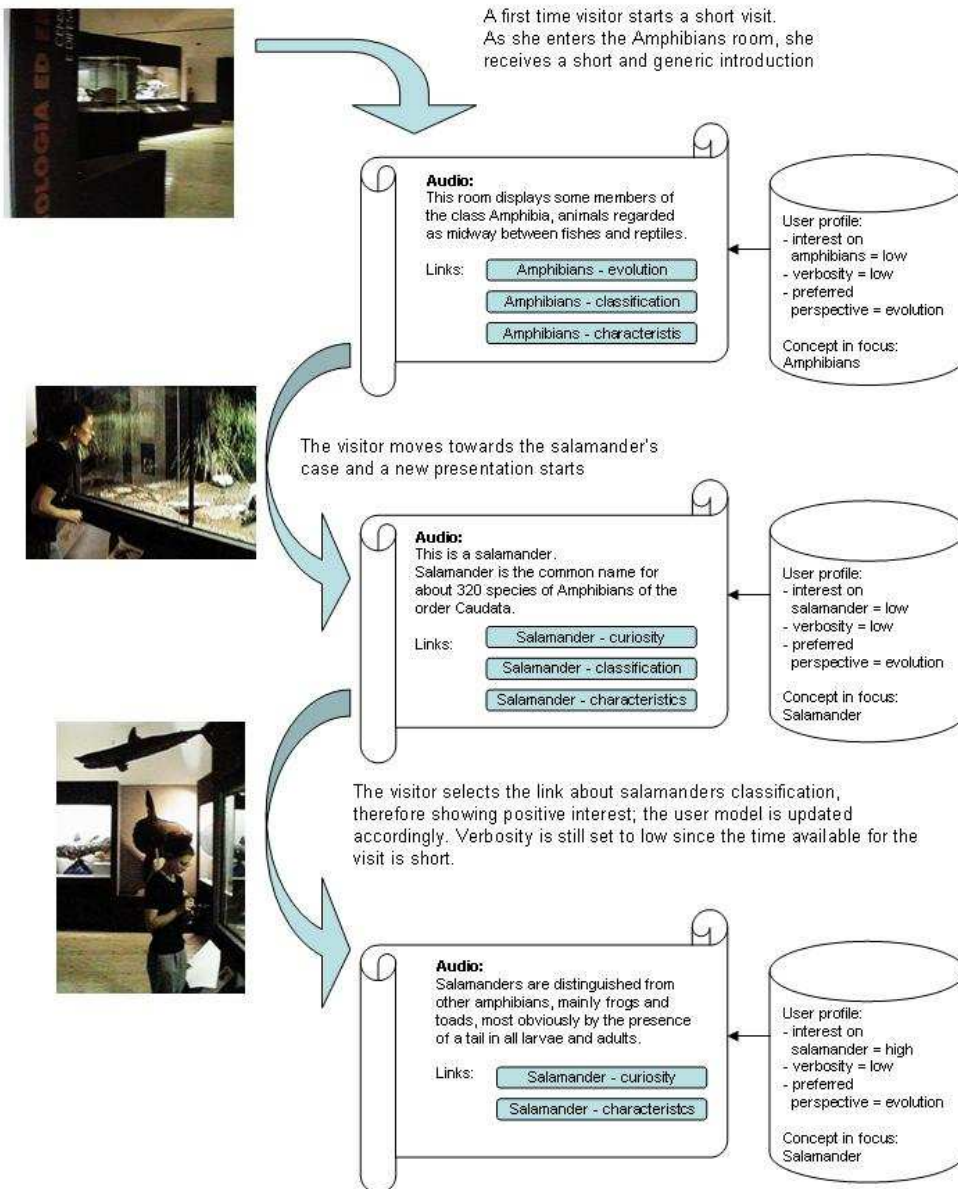
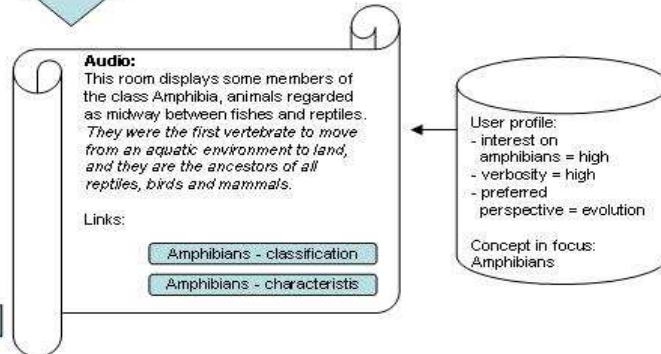


Figure 6. Scenario 1 for user interaction with HyperAudio.



A frequent visitor enters the Amphibians room after visiting the Reptiles room, where she has already seen the lizard. As a frequent visitor, she receives long and detailed descriptions. In the introductory presentation, for example, an elaboration about amphibians evolution (text in italics) is directly played to the user, instead of been showed as link as was in scenario 1.



The visitor moves towards the salamander's case. Given that the visitor has already seen the lizard, a comparison is made between the two animals, to support better understanding and learning. According to high verbosity, the presentation also directly includes an elaboration on amphibians evolution that for the visitor in scenario 1 was realized as a clickable link.

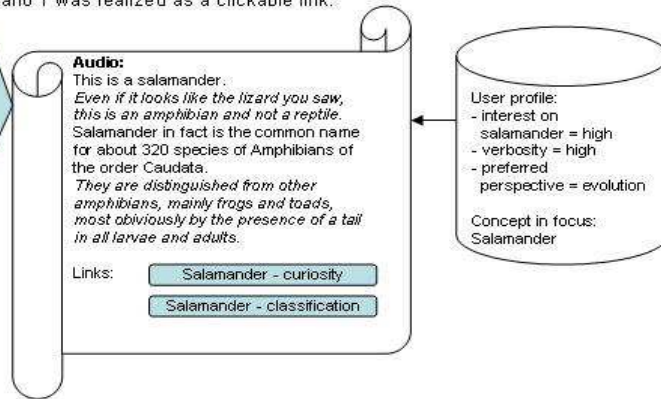


Figure 7: Scenario 2 for user interaction with HyperAudio.

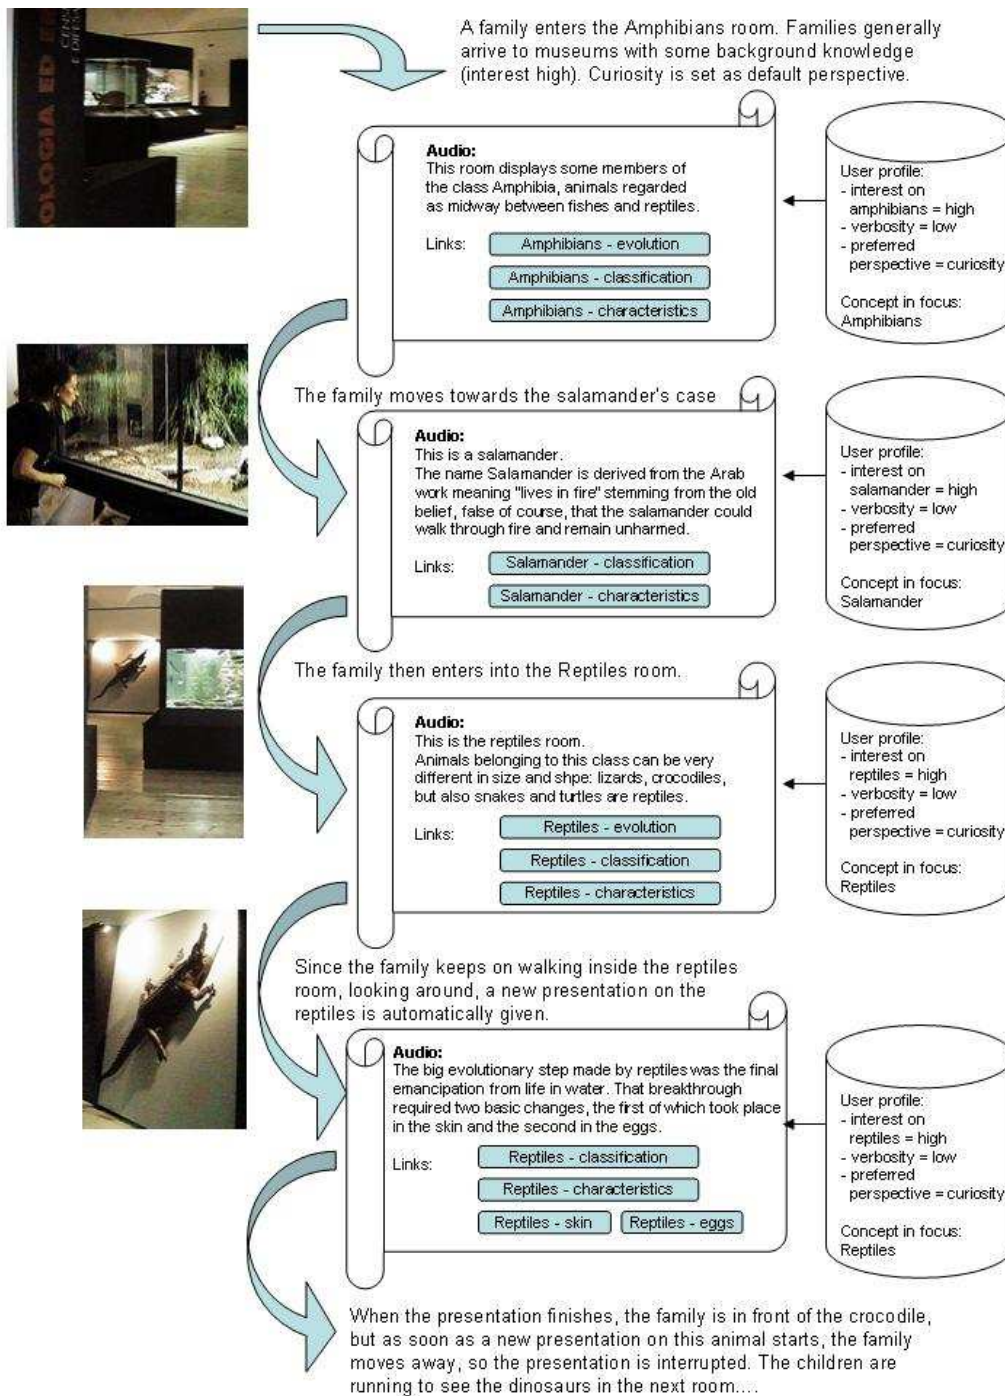


Figure 8: Scenario 3 for user interaction with HyperAudio

5. Rapid Prototyping and Testing

To assure a smooth interaction between the user and the final system, UCD advocates the application of an

iterative process of design, prototype, and test. Prototypes can be of different kind and are done for testing different system concepts; Houde and Hill (1997) have identified the following:

- a **role** prototype is used to test the function of the artefact in user's life;
- testing the **look and feel** means to concentrate on the interface and interaction;
- focussing on the **implementation** aspect would prototype techniques and components of the final system;
- finally, an **integration** prototype combines aspects of all the three prototypes above and moves the project towards its final form.

Although an integration prototype must be developed lastly, there is no particular order for the others that can be also done in parallel, e.g. testing the look-and-feel via paper mock-ups while testing how robust the localization is with an implementation prototype.

The more complex the system, the more important is to test each component separately before the integration step. Indeed a single-component test helps in focussing on only one aspect of the adaptive system. A multi-layer evaluation approach has been proposed for adaptive systems (Karagiannidis & Sampson 2000) and has demonstrated its power in localizing problems, i.e. in the interaction or else in the adaptation mechanisms (Brusilovski et al., 2001). Our approach is slightly more complex as we needed to take into account mobility as well as adaptivity.

When the implemented system is mobile and adaptive, testing the interaction in a multi-layer mode is more complicated as, for example, the evaluation should be delayed until the overall localization and communication infrastructure is fully functioning. However, an extensive testing of adaptive mechanisms can be done without involving the user, i.e. excluding the localization. To speed up the process of prototyping and testing, we developed an environment where components could be plugged in and tested while simulating others not-yet-ready modules. This approach was particularly useful in the HIPS project as different partners developed different components: using the development environment we were able to autonomously work on rules and data (the macronode network) while simulating the user model and the localization mechanism⁹. From that experience a set of guidelines can be proposed (Petrelli et al. 2000 reports the work in full, here only the most relevant points are discussed); Figure 9 visually summarizes the guidelines using HIPS as contextual example:

- **Modular architecture:** the adaptive hypermedia system and the development environment have to be designed simultaneously. The two architectures have to be similar if the environment has to support simulation of modules as well as component testing.
- **Plug and play:** adding or removing components should be easy. Plug-and-play finished components as well as easy disconnection of modules in need of the developer's attention is a very important feature when a team is involved. In HIPS during the tuning of macronodes and rules, the user model inquiry was simulated by manually setting different parameters; at a later stage the real modules were plugged in and the system was tested in full.
- **Component simulation via GUI:** to better support component simulation, a graphical interface should be offered for the easy setting of core values. This is particularly important if the values of the simulated modules are likely to change very often, as it was in our case for the location and orientation of the user.
- **Quick test-revise cycle:** since extensive testing is essential, setting a test should require just a few clicks. We have found useful a graphical panel for setting the hypothetical interaction context conditions, running the system and collecting the produced output.
- **Support localized testing:** besides manually setting the context values, it should be possible to manually set the data to be used in the test. The macronode to start with could be explicitly selected. This "localized testing" was very useful to discover specific problems in complex situations, e.g. how a specific adaptation rule works with a certain data configuration.
- **Cumulate the results:** it is useful to cumulate the output in a dedicated panel to support the monitoring of the behaviour of the system over time. This feature is essential for checking the

⁹ A quite elaborated knowledge and interest model (Oppermann & Specht, 2000) was developed at GMD and was later integrated with a dynamic visiting style model (Gabielli et al., 1999) developed at the University of Siena. The University of Siena was also responsible for the localization sub-system (Bianchi & Zancanaro, 1999).

adaptive system as a whole and to grasp what the user would experience while interacting.

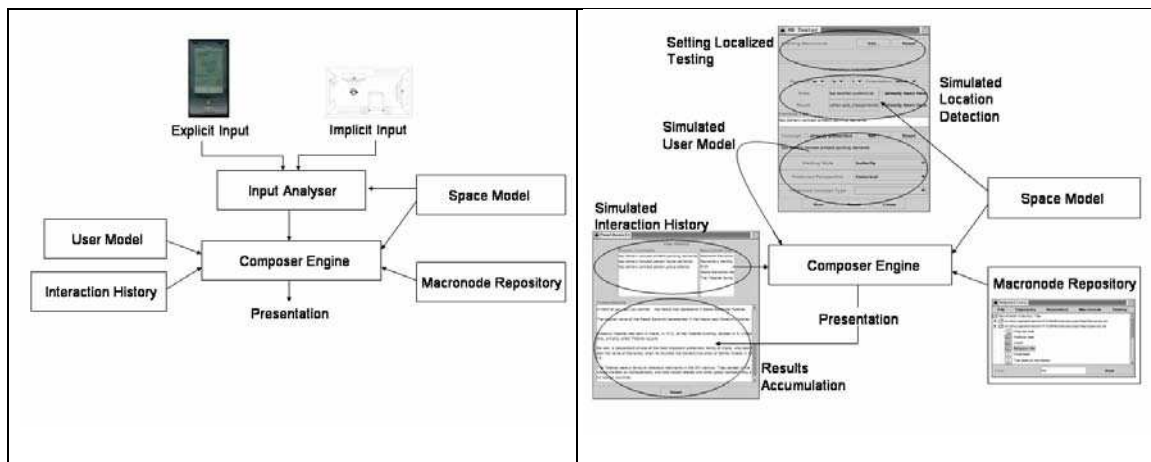


Figure 9. The system simplified architecture (left) compared to the development environment (right).

A development environment as outlined above allows considering and evaluating many context features at the same time as well as focussing on a single aspect of the complex adaptive mechanism (e.g. testing different user models while keeping the rest of the system fixed). The more complex the context the more valuable the help provided by the environment as all aspects of the context are related and can influence each other in negative and unwished ways. Consider a visitor listening to a presentation and moving toward an exhibit. Deciding if the presentation has to be stopped and a new one played may depend on factors other than the visitor's movements, for example the type of presentation currently playing: if it is about a specific object previously in focus it has to be interrupted (or better shortened to the end of the current sentence) as the references to the space are no longer valid, but if the content of the presentation is generic, then a full delivery is appropriate (the new presentation is queued). Using the environment it is easy to test different rules or combination of rules and evaluate the final effect thus shaping the adaptive behaviour precisely as the designer intended.

6. Content Editing and User Evaluation

A development environment that is to effectively support the overall UCD cycle for adaptive hypermedia cannot neglect the issue of content editing support. Content creation must be coupled with immediate testing and revision to guarantee coherent system behaviour. The macronode network used in the HyperAudio prototype was developed without any support and writing, connecting, and testing the 98 macronodes (for 7 objects, 2 rooms and 5 exhibits) proved to be error prone and time consuming. The cost of hand-writing was prohibitive for the bigger scale project HIPS and a editing support was deemed necessary.

The need for some editing support when authoring content for adaptive systems has been recognised in the past and recently addressed as an integral part of the development of adaptive hypermedia (Weber et al., 2001; De Bra et al., 2003). While the usefulness of graphical support has been recognized as complexity increases, the checking has only been considered at the level of graph consistence and rules propagation (Calvi & Cristea, 2002). Unfortunately this is not enough when linguistic adaptation is involved: checking for graph correctness would not say if a deictic reference was properly applied or the narration was fluent. A human has to systematically check the data structure and how the adaptation process uses it. The environment described in the previous section can be used to support an author in correctly creating content, by properly supporting the editing of the annotated network, and testing how the adaptive rules work on it. For one of the HIPS prototypes (the one used in the Museo Civico in Siena) a network of 170 macronodes was prepared to cover 31 exhibits in the museum; the total number of audio files created to support linguistic variation was 344. The same environment was later used in the M-PIRO project (Androutopoulos et al., 2001): 69 macronodes were created to cover 8 exhibits.

Using the environment has improved efficiency and effectiveness of adaptive hypertext editing. Lessons

were learnt, and suggestions can be given with specific reference on deploying an editing environment to adaptive hypertext authors external to the system development team, as it was done in the HIPS project. Figure 10 shows the components commented below and visually describes the relation between editing and testing.

- **Templates of (optimal) data organization:** developers should create templates of sub-networks that implement predefined directives to guide authors towards the correct compilation. The author can then concentrate on the content filling and the checking tasks. This feature is particularly important when the responsibility of creating the data is on the domain experts, i.e. museum curators. Through templates the developers can pass the basic knowledge on how the content had to be structured for optimal performance; by using examples of well-formed sub-networks authors can also gain a better understanding of the adaptive system.
- **Editing and testing the content network through a visually rich interface:** a basic display of the macronode network was available: the author could see at a glance the connections and the general content structure. Different views (by list or graph), searching facility, and user-defined data files turned out to be very handy features for network composed of hundreds of nodes. The possibility of getting at a glance an idea on, for example, the length of a presentation (i.e. length of a path in the network) or the type of content delivered (e.g. anecdotal or historical) was very useful for creating a balanced network where all the nodes got the chance to be selected and listened. A further improvement of the graphical interface is the progressive highlighting of the nodes used; this way the author can quickly check that all the nodes can be reached.
- **Quick test-revise cycle:** as for testing adaptivity, an extensive testing is essential in content editing as well. Simple test run, quick problem identification, and immediate fixing have to be supported. The features discussed in the previous section, namely a graphical panel for setting the hypothetical interaction context conditions, selecting the node to start with, running the system, and collecting the produced output, proved to be a valuable support for fast testing and problem identification. To support immediate fixing, the testing panels must be integrated with the editing ones so that the author can just turn her attention to the editing facilities for updating, switching then back for another test run.

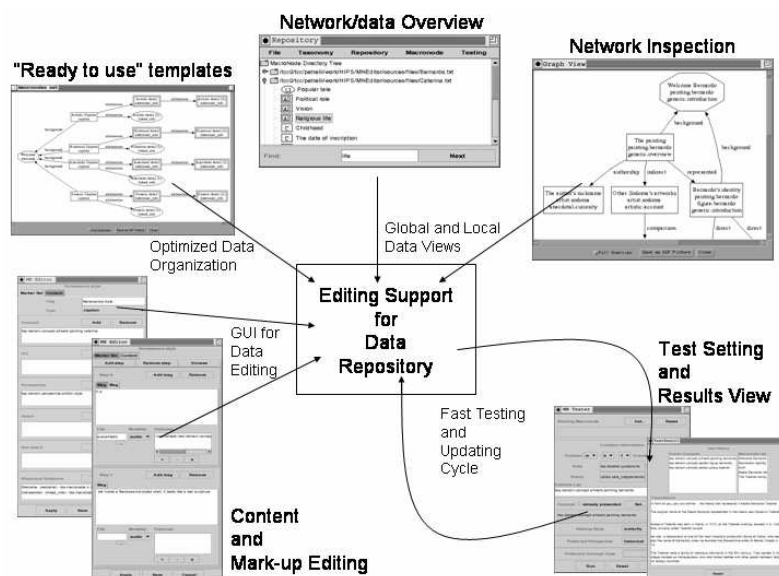


Figure 10. The editing supports components as related to the testing.

As discussed above, a comprehensive off-line/in-lab testing is essential to assure that the adaptive system is robust enough to go into users hands for a full evaluation. But this is just the first step. Indeed testing adaptive systems in real conditions is not trivial and the empirical evaluation of adaptive systems is a research

topic in itself (Weibelzahl & Paramitys, 2004). However as the number of adaptive systems implemented on mobile technology increases, there is the need of enlarging the context to location and device (as discussed by Gupta & Grover in Weibelzahl & Paramitys, 2004). New user evaluation methodologies are then needed to test adaptive and mobile guides, particularly when the use is not work-based but a leisure one (Marti, 2000; Marti & Lanzi, 2002). A first step in this direction can be found in this issue (Hatala & Wakkary, 2005): Hatala and Wakkary explore different dimensions of evaluating adaptive and mobile systems; they suggest, as we do, that an initial part of the evaluation (the “validation”) can be done off line (in the lab) to determine the most appropriate parameters to be used in, for example, the user model; the second phase, the “verification”, must be done with real users and should consider dimensions like variability, sustainability, and evolvability.

7. Related Work on Adaptive Mobile Systems

Adaptive mobile information systems for tourists and travellers is a popular topic¹⁰. It seems so promising that a generic user modelling system for tourist applications has been recently proposed (Fink & Kobsa, 2002). It is used by the WebGuide system to provide tourists in Heidelberg (Germany) personalized tours tailored to the user’s interests and preferences, transport facilities (e.g. car or bicycle), geographical distance, and specific user constraints (e.g. limited time) (Fink & Kobsa, 2002).

INTRIGUE (Ardissono et al., 2002) is another tour scheduler. It helps visitors in planning tours in Torino (Italy) and its surrounding area adapting to the needs of a group of people travelling together, e.g. parents and children. Users have to fill an initial “registration form” that provides the system with information on day and time of the visit, categories and geographical areas of interest: on the bases of this data INTRIGUE schedules a tour taking into account transfer time. It also adapts its layout according to the display device, desktop PC or WAP phone.

CRUMPET (Poslad et al., 2001) uses a handheld computer (a iPAQ) to provide personalized recommendations of services and attractions for city tourists, tour planning, proactive tips for nearby sites of potential interest, interactive maps and automatic adaptation to network services. Adaptation is based on a dynamic user interest model calculated from positive examples. The user can directly access and modify her user model. Stereotypes are mentioned as means for fast adaptation but it is not clear if the empirical studies alluded to for identifying the typical interest profile have ever been conducted (Poslad et al., 2001). The interface layout is simple and has been designed with computer-web literate users in mind.

GUIDE (Cheverst et al., 2000) implements a location-aware adaptive guide for tourists visiting the city of Lancaster (UK); it selects sites nearby the current user position, that are open, and are compatible with the user profile. It can also plan a tour of the city arranging sites selected by the user: the order depends on sites’ opening and visiting times, on distance and on the scenic route between sites. The user interface resembles a Web browser and new information is provided only after user interaction. Initially users are required to fill in a form asking for name, age, language, and interests. GUIDE also offers additional services such as booking accommodation, retrieving information (e.g. restaurants), and messaging with other tourists.

Hippie (Oppermann & Specht, 1999) is one of the preliminary prototypes developed inside the HIPS project. Conversely from the previously discussed applications, it is for inside use. It provides the visitors of an art exhibition with comments specific to the objects in sight; it adapts to user interests and knowledge calculated on the basis of actual behaviour. Hippie has a browser based interface that shows when new information is available by displaying a small blinking icon and playing a “earcon”: by clicking on the icon the new information is delivered as a hypermedia page with image, text and speech. Tours are generated and proposed to the user on the basis of her assumed interests. An initial setting of user’s interest profile is

¹⁰ See also Baus et al. (2004) for a selected critical comparison of map-based mobile guides offering services like helping users in orienteering in an unfamiliar city or accomplishing simple tasks, e.g. find a hotel nearby in Lol@, booking theatre tickets in Smartkom. Note, however, that none of the systems reviewed by Baus and colleagues seems to emphasize dynamic and adaptive content delivery: just a few can filter information on the bases of the user’s current position but no adaptation is applied. This can be due to the outdoor context where truly reliable localization and communication cannot be guaranteed: a correct localization and a continuous communication are mandatory for systems that aim at telling stories on the bases of the user’s movements. As a consequence, aspects related to narration and interactive environments have especially been exploited by projects devoted to inside use, particularly in museums and exhibitions.

available but not mandatory.

Besides having different domains, the applications above share the same idea of active users interacting with an adaptive guide within a browser paradigm: Initially users set their own profiles, later they can request adapted information (tours or descriptions) and access the result. However, with small devices, like those of PDA, the interface design is particularly critical and new interaction paradigms need to be explored, as noted by Cheverst (2002). HyperAudio attempted to overcome the limit of the screen and explored the idea of interacting with the space. Our system fuses adaptive information with the environment surrounding the user to create an adaptive immersive environment. Adaptation is done in respect to the user but also in respect to her actual position and current movement in the physical space, and is realized in terms of content selection, linguistic realization and appropriate synchronization. This idea was fully exploited inside HIPS, where a more sophisticated architecture was experimented for very fine linguistic adaptation. A better adaptation to the space and the narration was possible because of a new space model¹¹ and a deeper discourse context. The space model was finer grained (Bianchi & Zancanaro, 1999) thus allowing for deictic reference to near or far objects (“this is” vs. “in front of you”), as well as to objects located beside or behind the visitor. Similarly the macronode formalism was revised to support a richer discourse context for controlling the narration at the word level (Not & Zancanaro, 2000; Not & Zancanaro, 2001). A new way of modeling users solely on the basis of their movements was used in HIPS to adapt presentation length (Marti et al., 2001) while content selection used full models of user interest and knowledge (Opperman & Specht, 2000). Finally a new graphical interface to help users in locating artworks in the room by highlighting them on a 3D user-centred perspective reproduction of the room was implemented (Gabielli et al., 1999).

The concept of the disappearing computer has been extended by Zancanaro et al. (2003) who have enhanced the idea of adapted audio presentation built into HyperAudio and HIPS with a synchronized visual track for the described fresco: the pictures shown on the screen are animated via camera movements and shot transitions using cinematic techniques driven by the underlying content and rhetorical structure of the audio message (Zancanaro et al., 2003; Rocchi & Zancanaro, 2003; Callaway et al., 2005). The video-clips enhance the presentation, helping the visitors in locating described details in a complex and vast fresco, and demonstrate how computer technology can empower and enrich everyday activities, implementing the vision of augmented environments (CACM, 1993).

Monitoring user’s free movements for adapting presentations has inspired research in the area of wearable devices. In the system developed by Sparacino (2002) the user wears a private eye (a small transparent screen positioned in front of a single eye) where additional information for the object in view are displayed producing in this way a visual augmentation of the museum space. A Bayesian network is used to model both the user (interest and style -busy, selective, or greedy visitor-) and the appropriateness of the content (length and order). A set of video clips derived from 2 hours film on the exhibition represent the content. The video clip to be delivered is conditionally selected respect to the user model, the appropriate order of delivery, and its length. The selected video clip is displayed on the private eye with textual and pictorial details.

The LISTEN project (Zimmermann et al., 2003) instead explores the audio channel: the user carries only headphones and moves freely in an adaptive 3D-audio art museum. LISTEN merges technology developed in virtual reality (3D audio environments) and adaptive interfaces. Data mining techniques are used to model user interests, preferences and movements; the adaptation affects the presentation style (e.g. music, spoken text, and sound effects), the presentation content (e.g. facts, emotions, overview), length and volume. Clues for the user modelling are derived from the time, the position (of user and object), and the object of focus. Within LISTEN a unified framework for context-management was experimented to integrate the modelling of the user and the modelling of the context (Zimmerman et al., 2005).

¹¹ A fine-grained and robust space model is essential to build up a sophisticated content adaptation system. Indeed being able to model the user, the space, and the objects in the same system (as proposed by Carmichael et al. in this issue (Carmichael et al. 2005)) opens a spectrum of interesting new possibilities.

8. Conclusions

The examples discussed in the previous section show how adaptive hypermedia are branching out from the narrow path of adapting content and links to users sitting in front of a screen towards a broader way of adapting to the interaction context of users immersed in an augmented environment at a certain time and place. As scenarios of use for adaptive systems overcome the limit of desktop applications, system complexity will continue to increase. A robust methodology and appropriate development tools will increasingly be fundamental for successful designs particularly when mobile and ubiquitous computing is associated to adaptivity. HyperAudio has been one of the few adaptive projects where a user-centred approach was used to design the system and likely the first in the area of adaptive and ubiquitous guides. From that experience we have learnt how a deep understanding of users and uses is essential when designing adaptive systems to be used in highly-constrained conditions, as running efficiently on a PDA; in this context each design choice has to be evaluated and motivated. In this paper we have shown how the most sophisticated and advanced techniques could fail when compared with real use, and how a simpler solution can be equally effective. From our experience an effective design is based on few, relevant assumptions derived from actual user needs that spread on all the aspects of the adaptive system, i.e. user model, adaptive rules, and annotated data. Designers' creativity is then instrumental for deciding how flexibility should be implemented, i.e. which adaptive techniques can better support an effective and efficient use of the system. With this respect, the support of a dedicated environment is mandatory for an iterative development and testing of the final adaptive behaviour. In this way designers can explore and test different technical solutions and authors can be supported in the creation of the data.

Acknowledgements

Many people have been involved in the HyperAudio project along the years: Gregorio Convertino, Antonella De Angeli, Christian Girardi, Marcello Sarini, Oliviero Stock, Carlo Strapparava, and Massimo Zancanaro. We gratefully acknowledge everybody's contribution in thinking, designing and developing HyperAudio. We feel also indebted with the team of the HIPS European project for the many fruitful discussions and group work that helped our ideas and technologies keep on evolve. The HIPS European consortium included: University of Siena (coordinating partner), CB&J (France), GMD (Germany), ITC-Irst (Italy), SIETTE-Alcatel (Italy), SINTEF (Norway), University of Dublin (Ireland) and University of Edinburgh (Scotland).

A special thank you to the Civic Museum of Rovereto and to its director, Dr. Franco Finotti, who by showing us the museum and the (patented) hardware inspired the project.

Finally our gratitude to Nigel Ford for his encouraging comments on the paper, and to the reviewers and journal editors for the many fruitful and thorough suggestions.

References

- Abowd G., C. Atckeson, J. Hong, S. Long, R. Kooper, and M. Pinkerton: 1997, 'Cyberguide: A Mobile Context-aware Tour Guide'. *Wireless Network*, **3**, 421-433.
- Androutsopoulos, I., J. Calder, E. Not, F. Pianesi and M. Roussou: 2001, 'Multilingual Personalised Information Objects'. In: *Proceedings of the International Workshop on Information Presentation and Multimodal Dialogue*, Verona, Italy, pp. 25-29.
- Ardissono, L., A. Goy, G. Petrone, M. Segnan and P. Torasso: 2002, 'Ubiquitous User Assistance in a Tourist Information Server'. In: P. De Bra, P. Brusilovsky, R. Conejo (eds.): *Proceedings of the 2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Malaga, Spain, Springer:LNCS 2347, pp. 14-23.
- Baus, J., K. Cheverst, and C. Kray: 2004, 'A Survey of Map-based Mobile Guides'. In: L. Meng, A. Zipf, T. Reichenbacher (eds.): *Map-based Mobile Services: Theories, Methods and Implementations*, Springer Geosciences, Springer Verlag, Heidelberg, pp. 197-216.

- Bederson, B.:1995, 'Audio Augmented Reality: A Prototype automated tour guide'. *CHI95 (companion)*. Denver (Colorado-USA), pp. 210-211.
- Benelli G., A. Bianchi, P. Marti, E. Not and D. Sennati: 1999, 'HIPS: Hyper-Interaction within Physical Space'. In: *Proceedings of IEEE Multimedia Systems '99, International Conference on Multimedia Computing and Systems*, June 7-11, Florence, pp. 1075-1078.
- Benyon, D.:1993, 'Adaptive Systems: A Solution to Usability Problems'. *User Modeling and User-Adapted Interaction* **3** (1), 65-87.
- Bianchi, A. and M. Zancanaro: 1999, 'Tracking users' movements in an artistic physical space'. In: M. Caenepeel, D. Benyon and D. Smoth (eds.): *Community of the Future. Proceedings i³ Annual Conference*, pp. 103-106, http://www.i3net.org/ser_pub/publications/proceedings/i3ac99-proceedings.pdf (accessed 15.10.2004)
- Broadbent, J. and P. Marti: 1998, 'An experiential modality of exploring physical and virtual spaces'. In: *Proceedings of the i3 Annual Conference*, Nyborg, http://www.i3net.org/ser_pub/annualconf/abstracts/navigation/expmodality.html (accessed 1.11.2004).
- Broadbent, J., P. Marti, E. Not, R. Oppermann, and O. Rahlff: 1998, 'Report on user requirements and scenario description', HIPS: Esprit-25574 project, Restricted Report.
- Bontcheva, K.: 2001, 'The impact of Empirical Studies on the Design of an Adaptive Hypertext Generation System'. In: S.Reich, M. Tzagarakis and P. De Bra (eds.): *Hypermedia: Openness, Structural Awareness, and Adaptivity*, Springer: LNCS 2266, pp. 201-214.
- Borun, M. and J. Dritsas: 1997, 'Developing Family-Friendly Exhibits'. *Curator* **3**, 178-196.
- Brusilovsky, P.: 1996, 'Methods and Techniques of Adaptive Hypermedia'. *User Modeling and User-Adapted Interaction* **6**, 87-129.
- Brusilovsky, P.: 2001, 'Adaptive Hypermedia'. *User Modeling and User-Adapted Interaction* **11**, 87-110.
- Brusilovsky, P., C. Karagiannidis and D. Sampson: 2001, 'The Benefits of Layered Evaluation of Adaptive Applications and Services'. In: *Proceedings of the Workshop on Empirical Evaluations of Adaptive Systems*, UM01, (<http://art.ph-freiburg.de/um2001/proceedings.html> accessed 15.10.2004), pp. 1-8.
- CACM - *Communication of the ACM*: 1993, Computer Augmented Environments: Back to the Real World. 36 (7).
- Calvi L. and A. Cristea: 2002. 'Towards Generic Adaptive Systems: Analysis of a Case Study'. In: P. De Bra, P. Brusilovsky and R. Conejo (eds.): *Proceedings of the 2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Malaga, Spain, Springer:LNCS 2347, pp. 79-89.
- Callaway C., E. Not, A. Novello, C. Rocchi, O. Stock and M. Zancanaro: 2005, 'Automatic Cinematography and Multilingual NLG for Generating Video Documentaries'. *Artificial Intelligence*, in press.
- Carmichael, D. J., J. Kay and R. J. Kummerfeld: 2005 'Consistent modelling of users, devices and sensors in a ubiquitous computing environment' in this issue.
- Cheverst, K., N. Davies, K. Mitchell, and P. Smith: 2000, 'Providing Tailored (Context-Aware) Information to City Visitors'. In: *Proceedings of First International Conference on Adaptive Hypermedia and Adaptive*

Web-Based Systems, Trento, Italy, pp.73-85.

Cheverst, K., K. Mitchell and N. Davis: 2002, 'The Role of Adaptive Hypermedia in a Context-aware Tourist Guide'. *Communication of the ACM* **45**(5), 47-51.

Cheverst, K., H. E.Byun, D.Fitton, C.Sas, C.Kray, and N. Villar: 2005 'Exploring Issues of User Model Transparency and Proactive Behaviour in an Office Environment Control System.' In this issue.

Cristea A. and L. Aroyo: 2002, 'Adaptive Authoring of Adaptive Educational Hypermedia'. In: P. De Bra, P. Brusilovsky and R. Conejo (eds.): *Proceedings of the 2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Malaga, Spain, Springer:LNCS 2347, pp. 122-132.

De Bra, P., and L. Calvi: 1998, 'AHA! An open Adaptive Hypermedia Architecture'. *The New Review of Hypermedia and Multimedia* **4**, 115-139.

De Bra, P., A. Aqerts, B. Berden, B. de Lange, B. Rousseau, T. Santic, D. Smits, and N. Stash: 2003. 'AHA! The Adaptive Hypermedia Architecture'. In: *Proceedings of International Conference on Hypertext and Hypermedi*, Nottingham, UK, pp. 81-84.

Dieterich, H., U. Malinowsky, T. Kuhme and M. Schneider-Hufschmidt: 1993, 'State of the Art in Adaptive User Interaction'. In: M. Schneider-Hufschmidt, T. Kuhme and U. Malinowsky (eds.): *Adaptive User Interfaces: principles and practice*. Amsterdam, Elsevier Science Publishers, pp. 13-48.

Falk, J. and L. Dierking: 1992, *The Museum Experience*. Ann Arbor, MI, Whalesback Books.

Fink, J. and A. Kobsa: 2002, 'User Modeling for Personalized City Tours'. *Artificial Intelligence Review* **18**(1), 33-74.

Gabrielli, F., P. Marti, and L. Petroni : 1999, 'Environments as Interface'. In: M. Caenepeel, D. Benyon and D. Smoth (eds.): *Community of the Future. Proc. of the i3 annual conference*, Siena, Italy, pp. 44-47.

Hackos, J. T. and J.C. Redish: 1998, 'User and task analysis for interface design'. Wiley.

Hartson, H. R.: 1998. 'Human-computer interaction: Interdisciplinary roots and trends'. *The Journal of Systems and Software* **43**, 103-118.

Halliday, M.A.K. and R. Hasar: 1985, 'Language, context and text: Aspects of language in a social-semiotic perspective'. Deaking University Press.

Hatala, M. and R. Wakkary: 2005 'User Modeling and Semantic Technologies in support of a Tangible Interface', in this issue.

Houde, S. and C. Hill: 1997, 'What do prototypes prototype?'. In: M. Helander, T.K. Landauer, P. Prabhu (eds): *Handbook of Human-Computer Interaction* (2nd edition). Elsevier Science, pp. 367-381.

Jameson, A. and E. Schwarzkopf: 2002, 'Pros and Cons of Controllability: An Empirical Study'. In: *Adaptive Hypermedia and Adaptive Web-Based Systems*, 2nd International Conference, AH 2002, Malaga, Spain, Springer: Lecture Notes in Computer Science 2347, pp. 193-202.

Johnsons, C.: 1998, 'First Workshop on Human-Computer Interaction with Mobile Devices'. <http://www.dcs.gla.ac.uk/~johnson/papers/mobile/HCIMD1.html> (accessed 20.8.2004).

Karagiannidis, C. and D. Sampson: 2000, 'Layered Evaluation of Adaptive Applications and Services'. In: P. Brusilovsky, O. Stock and C. Strapparava (eds.): *1st Int. Conf on Adaptive Hypermedia and Adaptive Web-*

Based Systems, AH 2000, Trento, Italy, Springer: LNCS 1892, pp. 343-346.

Kobsa, A., J. Koenemann and W. Pohl: 2001, 'Personalised Hypermedia Presentation Techniques for Improving Online Customer Relationships'. *The Knowledge Engineering Review* **16**(2), 111-155.

Lozowski Boisvert, D. and B. Jochums Slez: 1995, 'The relationship between exhibit characteristics and learning-associated behaviour in a science museum discovery space'. *Science education* **79**(5), 503-518.

Mann, W.C. and S. Thompson: 1988, 'Rhetorical Structure Theory: Toward a functional theory of text organization'. *Text* **8**(3), 243-281.

Marti, P.:2000, 'Report on The User Evaluation'. HIPS: Esprit-25574 project, Restricted Report

Marti, P., F. Gabrielli and F. Pucci: 2001, 'Situating Interaction in Art'. *Personal and Ubiquitous Computing* **5**, 71-74.

Marti, P. and P. Lanzi: 2001, '"I enjoyed that this much!" techniques for measuring usability in leisure-oriented applications'. In: Bawa, Dorazio, Trenner (eds.): *The Usability Business: Making the Web Work*. Springer-Verlag, pp. 49-59.

Milosavljevic, M., A. Tulloch and R. Dale: 1996, 'Text generation in a dynamic hypertext environment'. In: *Proceedings of the 19th Australasian Computer Science Conference*, Melbourne, Australia, pp. 417-426.

Milosavljevic, M., P. Brusilovsky, J.D. Moore, J. Oberlander and O. Stock (eds.): 1997, 'Proceedings of the Flexible Hypertext Workshop', held at the Eight ACM International Hypertext Conference, Southampton, UK. Macquarie University Technical Report C/TR97-06.

Nielsen, J.: 1993, 'Usability engineering'. New York: Academic Press.

Norman, D.: 1998, 'The Invisible Computer – Why good product can fail, the personal computer is so complex, and information appliances are the solution'. MIT press.

Not, E., D. Petrelli, O. Stock, C. Strapparava and M. Zancanaro: 1997a, 'Augmented Space: Bringing the Physical Dimension into Play'. In: *Proceedings of Flexible Hypertext Workshop* held in conjunction with Hypertext97, Southampton, UK.

Not, E., D. Petrelli, O. Stock, C. Strapparava and M. Zancanaro: 1997b, 'Person-oriented Guided Visits in a Physical Museum'. In: D. Bearman and J. Trant (eds.): *Museums Interactive Multimedia 1997: Cultural Heritage Systems. Design and Interfaces*, (Selected papers from ICHIM97), Archives & Museum Informatics, pp. 69-79.

Not E., D. Petrelli, M. Sarini, O. Stock, C. Strapparava, and M. Zancanaro: 1998, 'Hypernavigation in the Physical Space: Adapting Presentations to the User and to the Situational Context'. *The New Review of Hypermedia and Multimedia Journal* **4**, 33-45.

Not, E. and M. Zancanaro: 1998, 'Content Adaptation for Audio-based Hypertexts in Physical Environments'. In: *Proceedings of the 2nd Workshop on Adaptive Hypertext and Hypermedia*, held in conjunction with Hypertext'98, <http://www.wis.win.tue.nl/ah98/Proceedings.html> (accessed 4.11.2004).

Not, E. and M. Zancanaro: 2000, 'The MacroNode Approach: Mediating Between Adaptive and Dynamic Hypermedia'. In: *Proceedings of the 1st International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Trento, Italy, pp 167-178.

Not, E. and M. Zancanaro: 2001, 'Building Adaptive Information Presentations from Existing Information

Repositories'. In: *Proceedings of the International Workshop on Information Presentation and Multimodal Dialogue*, Verona, Italy, pp. 95-99.

Oberlander, J., M. O'Donnell, A. Knott and C. Mellish: 1998. 'Conversation in the museum: experiments in dynamic hypermedia with the intelligent labelling explorer'. *The New Review of Hypermedia and Multimedia Journal*, **4**, 11-32.

Oppermann R. and M. Specht: 1999, 'A Nomadic Information System for Adaptive Exhibition Guidance'. *Archives and Museum Informatics* **13**, 127-138.

Oppermann, R. and M. Specht: 2000, 'A Context-sensitive Nomadic Information System as an Exhibition Guide'. In: *Proceedings of the 2nd International Symposium on Handheld and Ubiquitous Computing*, HUC 2000, pp. 127-142.

Petrelli, D., A. De Angeli and G. Convertino: 1999, 'A User-centred Approach to User Modeling'. In: *Proceedings of the 7th International Conference on User Modeling*, Banff, Canada, pp. 255-264.

Petrelli, D., E. Not, M. Zancanaro, C. Strapparava and O. Stock: 2001, 'Modelling and Adapting to Context'. *Personal and Ubiquitous Computing* **5**, 20-24.

Petrelli, D., D. Baggio and G. Pezzulo: 2000, 'Adaptive Hypertext Design Environments: Putting Principles into Practice'. In: P. Brusilovsky, O. Stock and C. Strapparava (eds.): *Proceedings of the 1st International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, AH 2000, Trento, Italy, Springer: LNCS 1892, pp. 202-213.

Preece, J, Y. Rogers and H. Sharp: 2002, 'Interaction design: Beyond human-computer interaction'. Wiley.

Poslad, S., H. Laamanen, R. Malaka, A. Nick, P. Buckle and A. Zipf: 2001, 'CRUMPET: Creation of User-friendly Mobile Services Personalized for Tourism'. In: *Proceedings of the 2nd International Conference on 3G Mobile Communication Technologies*, London, pp. 28-32.

Rich, E.:1989, 'Stereotypes and user modeling'. In: A. Kobsa and W. Wahlster (eds.): *User models in dialog systems*. Berlin: Springer, pp. 35-51.

Rocchi C. and M. Zancanaro: 2003, 'Generation of Video Documentaries from Discourse Structures'. In: *Proceedings of the 9th European Workshop on Natural Language Generation*, Budapest, pp. 95-102.

Sarini, M. and C. Strapparava: 1998, 'Building a User Model for a Museum Exploration and Information-Providing Adaptive System'. In: *Proceedings of the 2nd Workshop on Adaptive Hypertext and Hypermedia, held in conjunction with HYPERTEXT'98*, Pittsburgh, USA, <http://www.wis.win.tue.nl/ah98/Sarini/Sarini.html>, (accessed 23.2.2005).

Sparacino, F.:2002, 'The Museum Wearable: real-time sensor-driven understanding of visitors' interests for personalized visually-augmented museum experiences'. In: *Proceedings of Museums and the Web (MW 2002)*, Boston, http://xenia.media.mit.edu/~flavia/Papers/flavia_mw2002.pdf (accessed 23.2.2005).

Vassileva, J.: 1996, 'A Task-Centered Approach for User Modeling in a Hypermedia Office Documentation System'. *User Modeling and User-Adapted Interaction* **6**, 185-223.

Weber G., H-C. Kuhl and S. Weibelzahl: 2001, 'Developing Adaptive Internet Based Courses with the Authoring System NetCoach'. In: S. Reich, M. Tzagarakis M. and P. De Bra (eds.): *Hypermedia: Openness, Structural Awerness, and Adaptivity*. Springer, pp. 226-238.

Weibelzahl, S. and A. Paramitys: 2004, 'Proceedings of the 3rd Workshop on empirical Evaluation of

Adaptive Systems'. In conjunction with AH2004

<http://www.easy-hub.org/hub/workshops/ah2004/proceedings.html> (accessed 1.11.2004).

Zancanaro, M., O. Stock and I. Alfaro: 2003, 'Using Cinematic Techniques in a Multimedia Museum Guide'. *Proceedings of Museum and the Web*, Charlotte, USA, pp. 209-216.

Zimmermann A., A. Lorenz and M. Specht: 2003, 'User Modeling in Adaptive Audio-Augmented Museum Environments'. In: *Proceedings of the 9th International Conference on User Modeling*, UM2003, Springer: LNAI 2702, pp. 403-407.

Zimmermann, A., M. Specht and A. Lorenz: 2005 'Personalization and Context-Management' in this issue.

Authors' Vitae

Daniela Petrelli

Department of Information Studies, University of Sheffield, Regent Court, 211 Portobello Street, S1 4DP Sheffield, UK

Daniela Petrelli is a lecturer (i.e. a UK faculty member status [maybe equivalent to US assistant professor]) in the Department of Information Studies at the University of Sheffield (UK) since 2001. Previously she was a Senior Research Scientist at ITC-irst in Italy where the work published in this volume was carried out. Dr. Petrelli received a Doctorship in Information Science from the University of Turin (Italy) and has a Diploma in Fine Art. Her research interests include exploiting ideas of User Centred Design in non-traditional areas of computer use, i.e. multimodal interaction, intelligent interfaces, multilingual information access, Semantic Web editing and interaction. The work published in this issue on the user centred design of mobile and adaptive museum guides fully represents her activity.

Elena Not

Cognitive and Communication Technologies Division, ITC-irst, via Sommarive, I-38050 Trento, Italy

Elena Not received a degree in Computer Science from the University of Udine (Italy). Since 1992 she has been a Research Scientist at the Cognitive and Communication Technologies Division at ITC-irst. She worked on several national and international research projects reflecting her interest in natural language generation with particular focus on text planning, discourse structure, referring expressions generation and disambiguation, dynamic hypermedia generation. The joint research with Daniela Petrelli described in this volume reflects her substantial experience in the development of adaptive multimedia information presentation systems.