User Modeling and User-Adapted Interaction (2005) 15:303–338
DOI 10.1007/s11257-005-8816-1

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

DANIELA PETRELLI[1,3] and ELENA NOT[2]
[1]*Department of Information Studies, University of Sheffield, Regent Court, 211 Portobello street, S1 4DP Sheffield, UK. e-mail: d.petrelli@shef.ac.uk*
[2]*Cognitive and Communication Technology Division, ITC–irst, Italy. e-mail: not@itc.it*

(Received: 21 April 2004; accepted in revised form 8 November 2004)

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 a system based on what is actually needed. This paper discusses the case of HyperAudio, a context-sensitive adaptive and mobile museum guide developed in the late 1990s. 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 define the 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. This is 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 in an iterative design process that considers the user interaction to be the central point. This paper discusses how such an environment allows designers and developers to experiment with different system behaviours and to widely test it under realistic conditions by simulating the actual context evolving over time. The understanding gained in HyperAudio is then considered from the perspective of later developments: our findings still appers to be valid despite the time that had passed.

Key words. Content adaptation, development support environments, flexible hypermedia, mobile guides, user centred design.

1. Introduction

To guide the design of information systems towards actual user needs and expectations, 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 evaluations of incremental prototypes (Preece et al., 2002). However, a rigorous user-centred design does not start with a prototype, but

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

with an extensive analysis of potential users, tasks, and environment (Hackos and Redish, 1998). Multiple techniques can be used, and the analysis of the data collected 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 (Harston, 1998; Nielsen, 1993).

The 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 and sometimes the interface layout (e.g., Bontcheva, 2001; Vassileva, 1996), but a pervasive user-centred design has hardly ever influenced the information organization or adaptation rules. Instead, a deep understanding of the 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 and where 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, see Kobsa et al., 2001). Choosing the required complexity is a design decision, which should be based on the knowledge that was acquired about the users and their tasks during the preliminary studies.

However, a good start is not enough to assure that the final adaptive systems will 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 will be 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 performed 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, . . .). 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 has it received the necessary attention from both practical and theoretical perspectives

(Cristea and Aroyo, 2002; Calvi and Cristea, 2002; De Bra et al., 2003; Weber et al., 2001). Still, data creation and rules testing are kept apart, possibly because content creation is considered the task of domain experts while rule testing is the 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 to test 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, e.g. user preferences, knowledge and behaviour, and interaction context.

When, in the mid 1990s, we started working on one of the first prototypes of an 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 on the application of UCD to adaptive systems. In the initial critical phase of the project we faced problems such as envisaging credible scenarios of use, identifying parameters for adaptivity, and designing content and adaptation rules in a suitable manner. The initial aim we had in mind was to offer the visitor personalized information centred on his/her current location. 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 free movement in an information space, and of automatically receiving tailored information. We started with the idea of adaptive hypermedia displayed on a PDA for browsing, and ended up with an intelligent system that took social and relational conditions, of the pace of the visit, and visitor's interests into account. It was intended to be a guide; it ended being a companion.

This profound change in how the adaptation should manifest itself was due to an extensive survey of museum visitors coupled with an explorative design process, as explained in the rest of this paper. Our analysis of visitors was not limited to descriptive statistics (e.g. the percentage of visitors who arrive at the museum already well informed) but also included correlating different data (e.g. of those likely to be families) and ultimately designing solutions (e.g. to 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 Hypernavigation In the Physical Space (HIPS; Benelli et al., 1999), a broader European project funded in the Intelligent Interactive Interfaces (i3) framework, where we further explored the UCD approach by creating a workbench for fast prototyping and off-line testing. The use of such a development environment tremendously facilitated the application of UCD to adaptive systems: we could test different solutions by a simple "plug-and-play" of 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's 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, we consider our insight into the evolution of design still valuable and unique after all these years. This paper reports on the HyperAudio experience of applying the UCD approach in the development of 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 regarding the 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 on immersive and adaptive mobile guides.

2. HyperAudio: Location Awareness and Adaptivity

2.1. THE HISTORY

The late 1990s saw a substantial increase in research on adaptive hypermedia in 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. Scientists in natural language processing were developing dynamic hypertext, where pages are generated on the fly on the basis of some domain knowledge representation, as well as a user model (Milosavljevic et al., 1996; Oberlander et al., 1998). 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 1990s (Wellner et al., 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 (Abowd et al., 1997; Bederson, 1995).

Our project started in 1997 with the aim of fusing these hot topics 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 museum visitors move in the physical space looking for interesting exhibits and wishing to acquire information to deepen their knowledge and satisfy their interests. The museum was envisioned as a sort of augmented environment, sensitive to visitors' movements, where an information hyperspace can be associated with each exhibit. Visitors would explore that hyperspace during the physical visit (Not et al., 1997a,b). The envisaged system would automatically play an audio commentary 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 each individual user, taking into account not only their interaction with the system, but also the broad interaction context, including the physical space, the visit so far, the interaction history, and the presented 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 the visitor's goal of enjoying the museum, and to ensure that visitors would find the visit 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 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 it had to consider what the visitor had already seen (in the physical space) or heard (from the hypermedia space). A properly designed adaptive guide would not propose the same information again to a visitor who returns 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, and hence the composition of the presentation had to consider rules for effective content structuring and linguistic realization according to the current discourse context. For example, 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 and 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, as in the use of the word "conversely" in the example below:

> "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 ahead of them are detected. Each meaningful physical location (e.g., exhibit, door, or passage) has a small power-autonomous infrared emitter that continuously sends out a unique code. Thus the physical space is partitioned into sensitive zones that allow the system to identify the visitor's position and orientation (the *Space Model* in Figure 2).

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. Pointing the pen on a displayed link (*explicit input*) activates the system as well, as outlined in Figure 2.

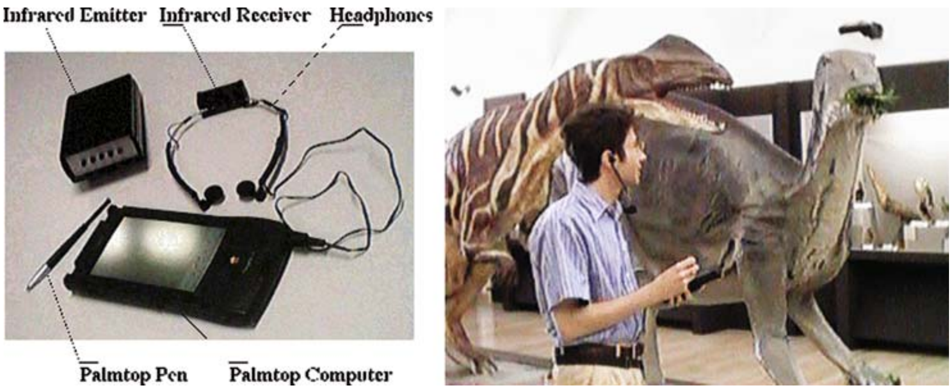

Figure 1. The Hyperaudio hardware.

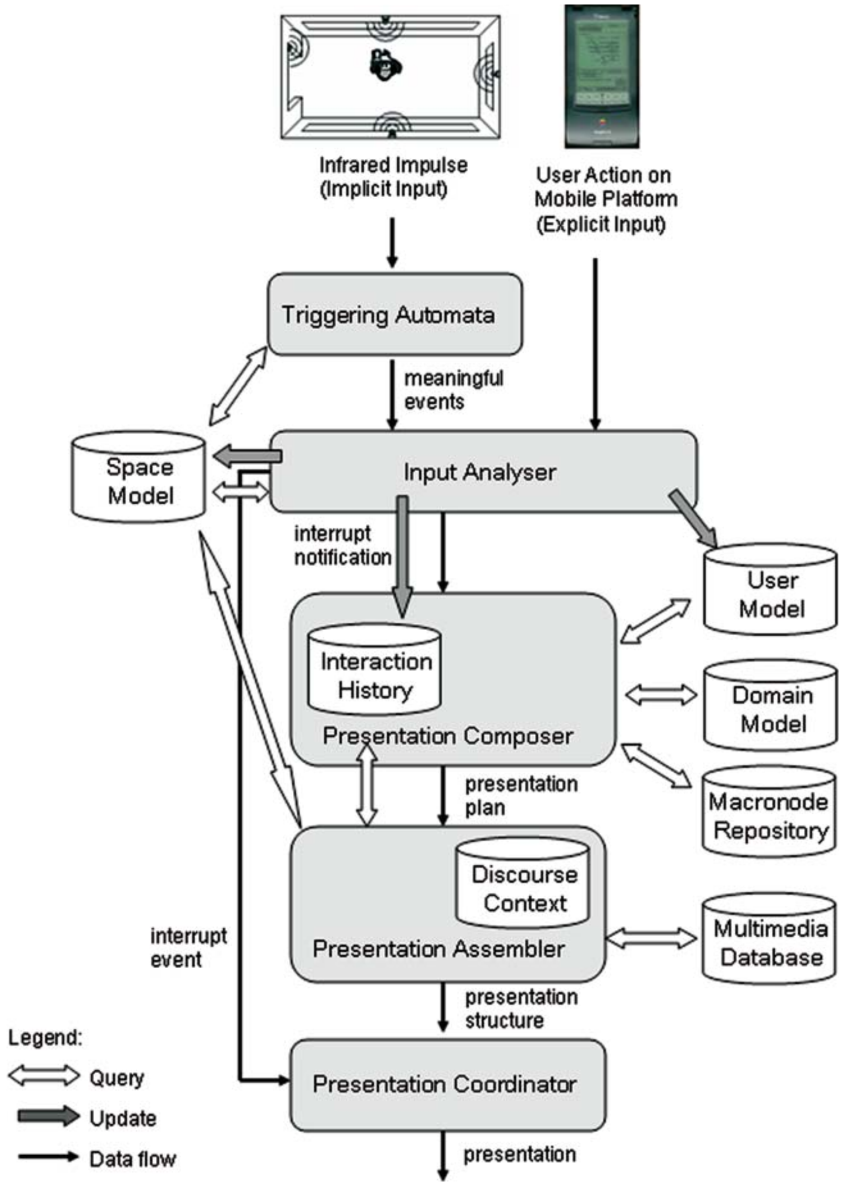

Figure 2. The Hyperaudio software architecture.

While the selection of a link at 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., if the visitor significantly changed her position), and asks the *Presentation Composer* for a *new presentation* that is appropriate for the new position. It also updates the *User Model* following the visitor's actions; for example, stopping a presentation shows evidence of a lack of interest 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 should be included in the audio presentation, which should become links and which ones should 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 assembles the final message. Here is where the linguistic arrangement takes place with respect to the current *discourse context* (see Section 2.3.2 for more information). 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 is 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 and Calvi, 1998). Dynamic hypermedia, instead, 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: it possesses 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 and Zancanaro, 1998). Strategies borrowed from the field of natural language generation are used to select and structure information units and to properly assemble multimedia pages (where audio plays a major role), thereby adjusting the linguisticrealization of the message to guarantee the coherence and cohesion of the final presentation.

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 and Thompson, 1988) that help describe the semantic relations between the various information units and how they may be textually integrated in a coherent manner. A macronode is internally organized to allow for some linguistic variation. Figure 3 shows a sample fragment of a 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 on the space model, the discourse context and the interaction history.

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",

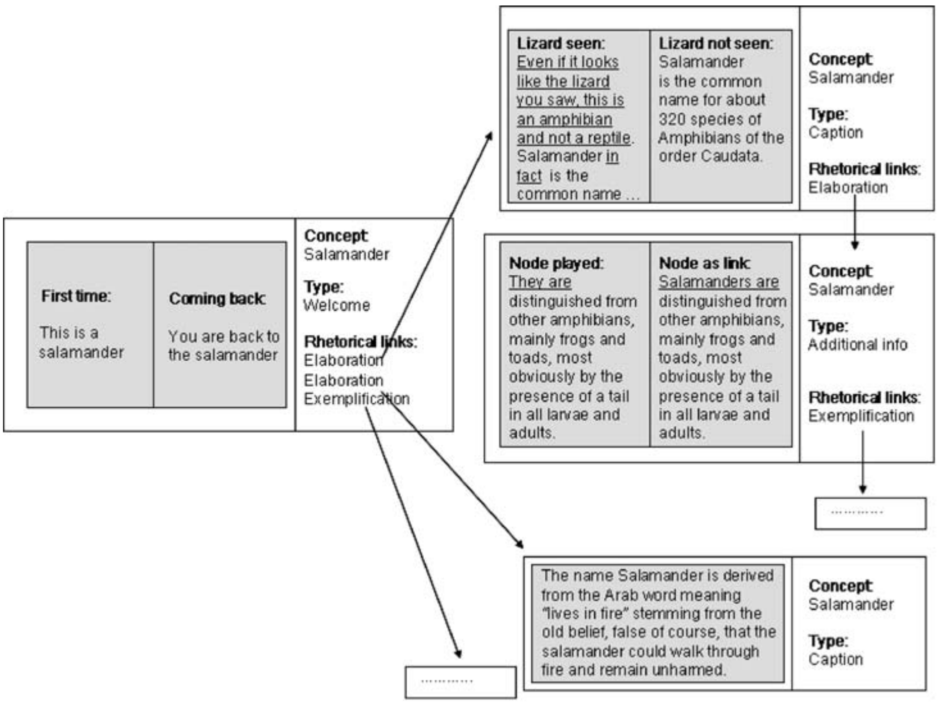

Figure 3. A network fragment of (simplified) macronodes: linguistic variations are underlined in the content part (left).

- "it is a cold-blooded animal",
- "this salamander is a cold-blooded animal".

In the original implementation, all linguistic adjustments of the surface form of macronodes were realized through conditional text that was 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. More recent research (Not and Zancanaro, 2001) 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.

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

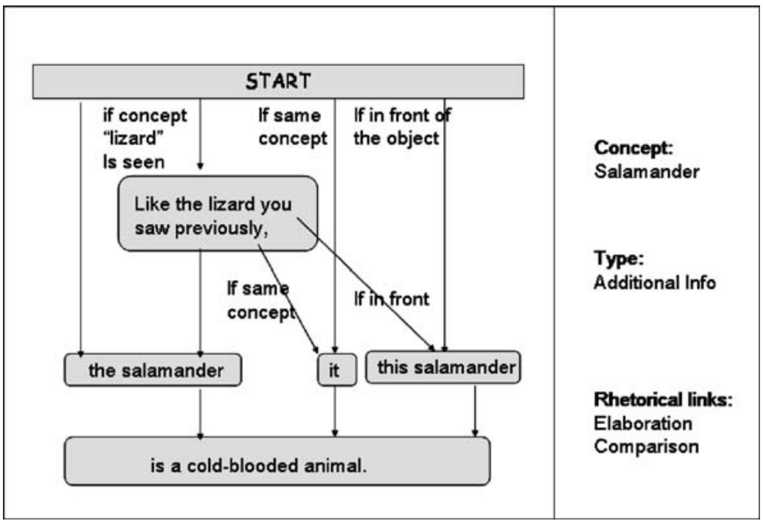

Figure 4. The internal structure of a macronode.

information for which the user's interest is high, and to present new information when the user goes back to a previous topic. In addition, rules checking the rhetorical links between macronodes control, for example, the length of an elaboration chain and the inclusion of background information to clarify a topic.

The Presentation Assembler takes care of adjusting the linguistic form of the presentation to 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') " or "if a concept has been already introduced (e.g., the object has been seen), then include an explicit reminder (e.g., 'you previously saw') ".

2.4. USER INTERFACE

The design of the interface was based on two fundamental constraints: (i) the MessagePad screen has a low resolution, and (ii) the visitor's attention is devoted to the exhibition and should not be unduly deflected. 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" function enables repeated listening to previously played presentations.

3. User Requirements Elicitation

As mentioned in Section 1, UCD typically starts with an extensive analysis of the potential users, the tasks and the 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 the adaptation rules. The main purpose was to identify the user characteristics that would compose the user profile. Implementing profiles as stereotypes (Rich, 1989) 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 many years and a whole museum literature is devoted to this topic. However, this extensive knowledge was of limited

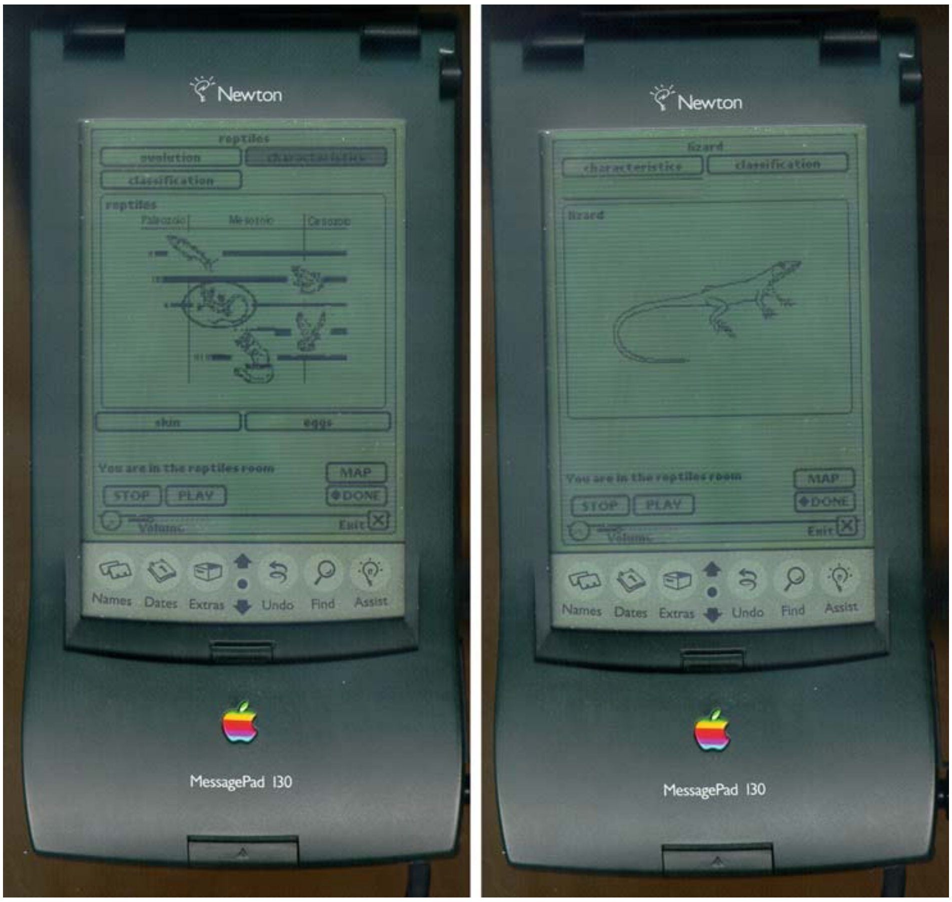

Figure 5. The interface layout as displayed (a) in the reptiles room, and (b) in front of the lizard.

help in defining stereotypes and adaptivity, since the focus is generally on how exhibition layout affects people's motion and how to make it effective in catching attention. 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. If confirmed, this would allow us to introduce explicit 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 only if they correlated with a predictable behaviour. The objective was to discover, for example, whether older people have a negative attitude toward technology and would prefer to be guided, in which case Hyperaudio would use non-interactive settings. Conversely, a positive attitude expected from younger people would be met by a highly interactive mode.

3.1. THE CASE STUDY

A survey was conducted to find out whether relations could be found between personal facts (e.g. age, specific interests) and the way museums are visited. A questionnaire was organized around five topics:

- A personal data profile section asked for demographic information, namely age, sex, education, job and place of residence, factors that we thought might have an impact on attitudes towards museums.
- A museum habit section complemented the personal profile. It collected data on how often the respondents visited museums and in what manner, e.g. alone, with a partner, with the family, or on a guided tour.
- A context section focused on the just completed visit and asked whether it was the visitor's first time in the museum, with whom they came, and about the general motivation for the visit.
- An itinerary section collected opinions on the use of guides (from books to human guides) as well as the duration and the purpose of the current visit.
- A styles of visit section collected general attitudes and opinions with regard to 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. 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 from those of younger ones; education was high for almost all museum visitors (91%); and 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 museum visitors for personal details at the beginning of the visit would not be useful. Fortunately, other attributes were discovered which accounted for visit and interaction variability.

Contrary to our prediction, the social dimension emerged as an important factor. Only 8% of visitors like to visit the museum alone; 24% prefer to be accompanied by friends and partners; 20% choose organized tours; and a majority of 42% go with the family. Visiting a Science 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 when discussion is indicated as being very important, 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, and 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 and Dritsas, 1997). Our data also showed that families are more likely than adult groups to arrive at the museum already informed. This indicated that classes of users had to be considered, with 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 correlated with the type and the duration of the visit. 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. For instance, 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 result in an autonomous self-sufficient style of visit.

To reinforce the previous finding, only 7% reported to like 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. Second, 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, and 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 what are the important aspects of visiting a natural science museum. From this knowledge a set of user requirements were extracted:

– A 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 distinct 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 context,[4] a system that autonomously decides 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 the finding of Cheverst et al. (2002) that during the evaluation of GUIDE, the vast majority of users wanted to invest as little effort as possible in navigation for information retrieval. The final HyperAudio prototype even supported a proactive mode that provides information fully automatically, thus requiring no interaction at all. Although a formal user evaluation never took place,[5] we observed many people using HyperAudio in a small museum simulation installed at ITC-irst: all

[4]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 and Schwarzkopf, 2002). For example, when a proactive adaptive system is used to support activities in a daily working environment, even allowing the user to scrutinize and modify the inner user model and system inference rules might be important (Cheverst et al., 2005).

[5]By the time the prototype was ready the MessagePad was no longer being produced or supported by Apple, thus the planned deployment in a museum setting never took place.

Table I. Preferred visit companion

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

Table II. Preferred visit support (not all options were included in both questionnaires; multiple choice was allowed in the HIPS case whose percentage figures have been correspondingly adjusted)

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

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,* i.e. 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 European HIPS project, a wider study of user requirements was conducted (Broadbert et al, 1998). The study was carried out at four different locations in three different countries (Norway, Germany and Italy), and focussed on art museums (modern art in Norway and Germany and historical palaces in Italy). Besides questionnaires distributed to visitors, focus groups with stakeholders (e.g. museum curators, art experts, custodians) were held to more precisely capture the needs of both visitors and managers. The goal of the questionnaire in this study was not precisely the same as in HyperAudio. The HIPS questionnaire focussed more on the art characteristics that visitors favour (to obtain empirical data on the design of an interest model) rather than on discovering actual behaviour and attitudes.

However, despite the differences in the two questionnaires some degree of comparison is possible.[6] The first noticeable difference is in the type of visitors (Table I) with a strong dominance of family and group in the science museums versus partner and friends in art museums. A second difference is in the preferred guide (Table II summarises the data). While human guides were the most preferred type of support in science museums, study participants definitely preferred a more autonomous visit in art museums. It is interesting that in both studies technological supports (the audio guide) were equally disliked.

[6]The original data is no longer available to perform the same analysis done for HyperAudio and to see if behaviour and attitudes in art museums differ from those in science museums.

Both tables clearly show how the two contexts (science and art) are different and underline how assumptions based on the research results of others can be risky. For example, caring particularly for family or group visitors does not seem justified in an art context. The need for a direct empirical analysis of the application domain is corroborated by other results from the HIPS study, which polled visitors on the many facets of interest in art (e.g. technique, composition, theme, artist, social or political context, history, etc.) and found two polarized interest clusters, namely historical vs. modern.

Although some results from previous studies can be generalized, imported, and re-used, such as the suggestion of including maps or guided tours in museum mobile guides (Broadbend et al., 1998; Broadbend and Marti, 1998), information that can influence adaptation needs to be collected anew by the new design team, and targeted toward the open questions that need direction.

4. System Redesign Based on User Requirements

Empirical evidence is used in UCD to direct system redesign and adjustment. In Hyper-Audio, this meant reconsidering the functionalities and adaptive behaviour the system was to support on the basis 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, rather than more complex reasoning. In the next few sections, we will explain how we revised the user model, the data structure and the matching rules with these considerations in or mind.

4.1. FROM A USER PROFILE TO A "VISIT" MODEL

The strongest effect of the user study was on the user model. Its original design was based on a thorough study of the existing literature on how visitors typically behave in museums (e.g. Falk and Dierking, 1992), conversations with museum curators, and studies on how exhibition layout affects visitors' behaviour (Lozowski and Jochums, 1995). We originally intended to maintain a fine-grained user profile whose data would be collected via an initial detailed questionnaire. The questionnaire data was intended to be merged with the predicted attracting and holding power of each exhibit[7] and used to initialize a model of the user's background knowledge, interest, and interaction preferences (Sarini and Strapparava, 1998).

As soon as the analysis of the user studies became available it was clear that some of our initial hypotheses about interest and knowledge modelling had to be revised. The idea of stereotypes based on personal features was abandoned and a

[7]Attracting power is the probability that the visitor stops and looks at an the exhibit. Holding power is the average time spent by visitors in front of it (Lozowski and Jochums, 1995).

more general shift of perspective from a user model to a *visit model* took place. In relation to this, four distinctions 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 choose an interaction mode (e.g., interactive for families, automatic for schools and adults).
- First-time visit: first-time and frequent visitors are different. 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 as well as the depth of descriptions.
- Interaction preferences: proactive behaviour is the default mode. However, it is considered 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 availability of new information is announced by 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 these questions would allow museum staff to enter the responses on behalf of the user when the guide is handed out, thus providing personalized information also to passive visitors, i.e. people who would never explicitly interact with the system.

The dynamic part of the user model was also revised. Initially conceived as a complex weighted activation network over domain concepts (Sarini and Strapparava, 1998), user interests were 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 a new presentation, those already heard are discarded. Finally a Boolean value to regulate system verbosity was introduced; it would be set to long visits or returning visitors. This very simple implementation of the dynamic user model had the advantage of being efficient even with limited computational power, as when using only a PDA. As a consequence, the response of HyperAudio to user movements was very fast and the natural pace of visits was not affected by the system.

4.2. REVISITING THE DATA FORMALISM

The macronode formalism discussed in Section 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. groups of adults) 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 is 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 Section 2.3.2, HyperAudio's adaptivity is realized by different sets of rules directed to different objectives. All rules were revised as a result of the requirements derived from our user study. Discovering that visitors might never interact suggested reinforcing the system reactivity to physical actions; for example "if the presentation has finished and the visitor does not move for the next two seconds, then prepare a new presentation".

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

New composition rules were also derived from the revision of the data structure described above. For example "if a user is a frequent visitor, than 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, namely "having fun" for the family, "learning the 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 scenarios in figures 6, 7 and 8 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.

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, prototyping, and testing. Prototypes can be of different kinds and are developed for testing different system concepts; Houde and Hill (1997) identified the following prototypes:

- a role prototype is used to test the function of the artefact in a user's life;
- testing the look and feel means to concentrate on the interface and interaction;
- for focussing on the implementation aspect one 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 at the end of the project there is no particular order for the others. They can be even 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 it 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 and Sampson, 2000) and has demonstrated its power in localizing problems, e.g. in the interaction or in the adaptation mechanisms (Brusilovski et al., 2001). Our approach is slightly more complex as we needed to not only take adaptivity into account but also mobility.

When the implemented system is mobile and adaptive, testing the interaction in a multi-layer mode becomes more complicated as, for example, the evaluation should be delayed until the overall localization and communication infrastructure

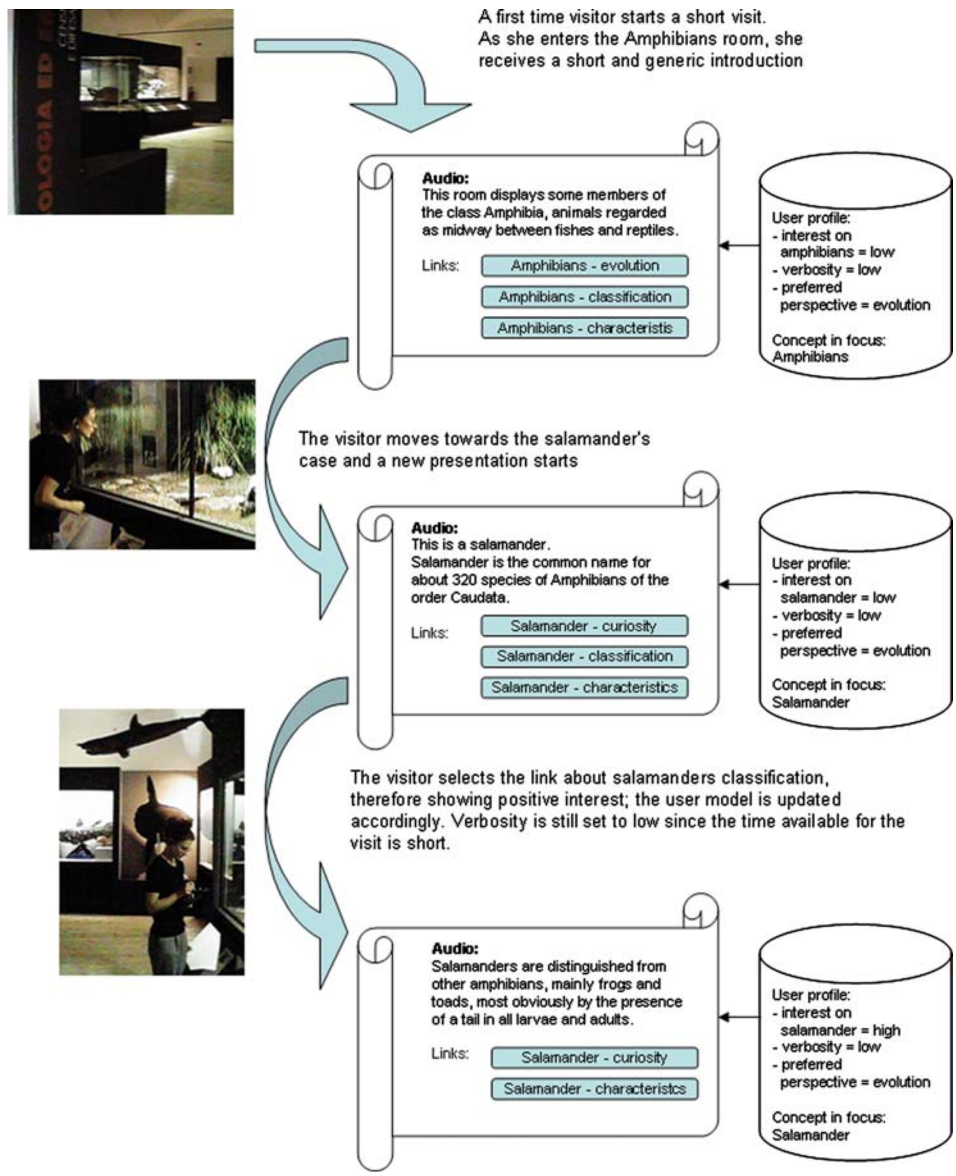

Figure 6. Scenario 1 for user interaction with Hyperaudio.

is fully functioning. However, extensive testing of adaptive mechanisms is still possible 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 the functionality of other modules that were not ready yet would be merely simulated. This approach was particularly useful in the HIPS project as different partners developed different components. With

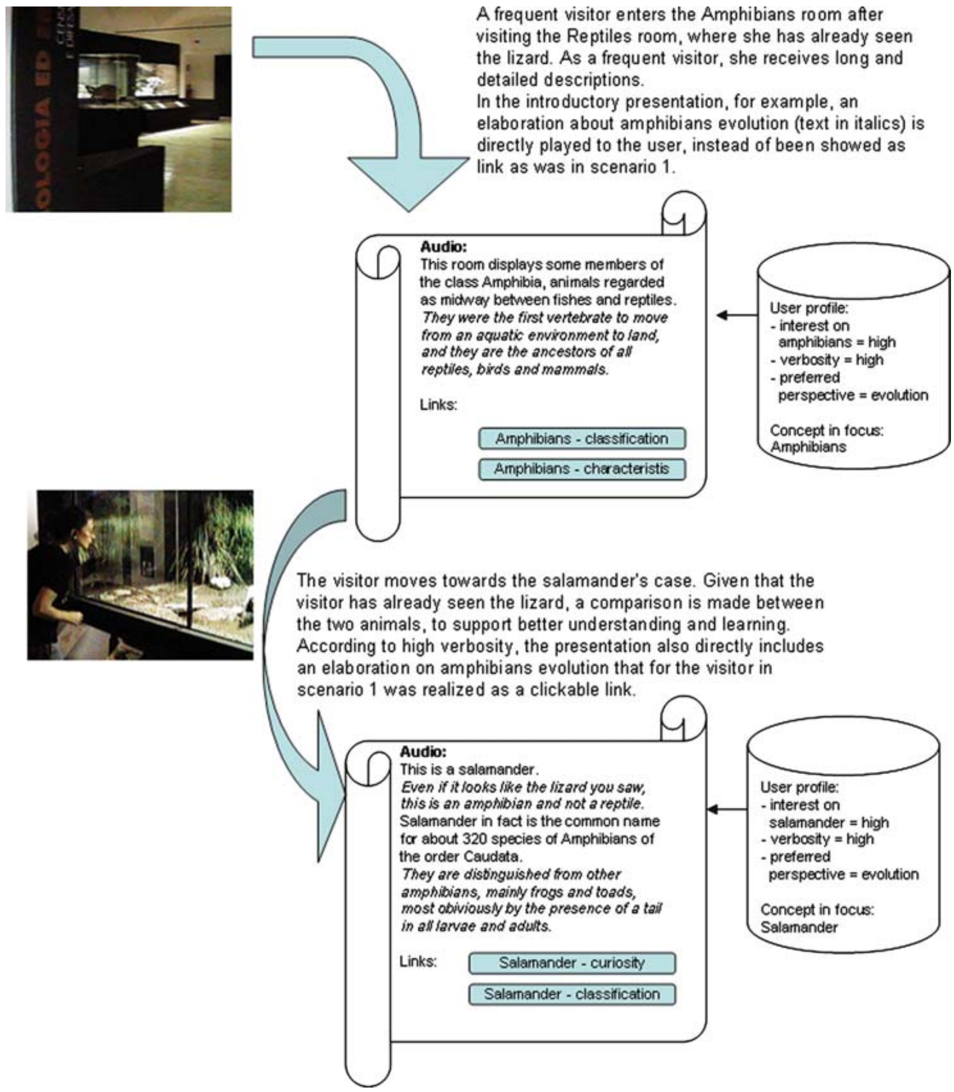

Figure 7. Scenario 2 for user interaction with Hyperaudio.

the development environment we were able to autonomously work on rules and data (the macronode network) while simulating the user model and the localization mechanism.[8] Based on these experiences, Petrelli et al. (2000) proposed a set of design guidelines. Figure 9 visually summarizes the most relevant of them, using HIPS as a contextual example:

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

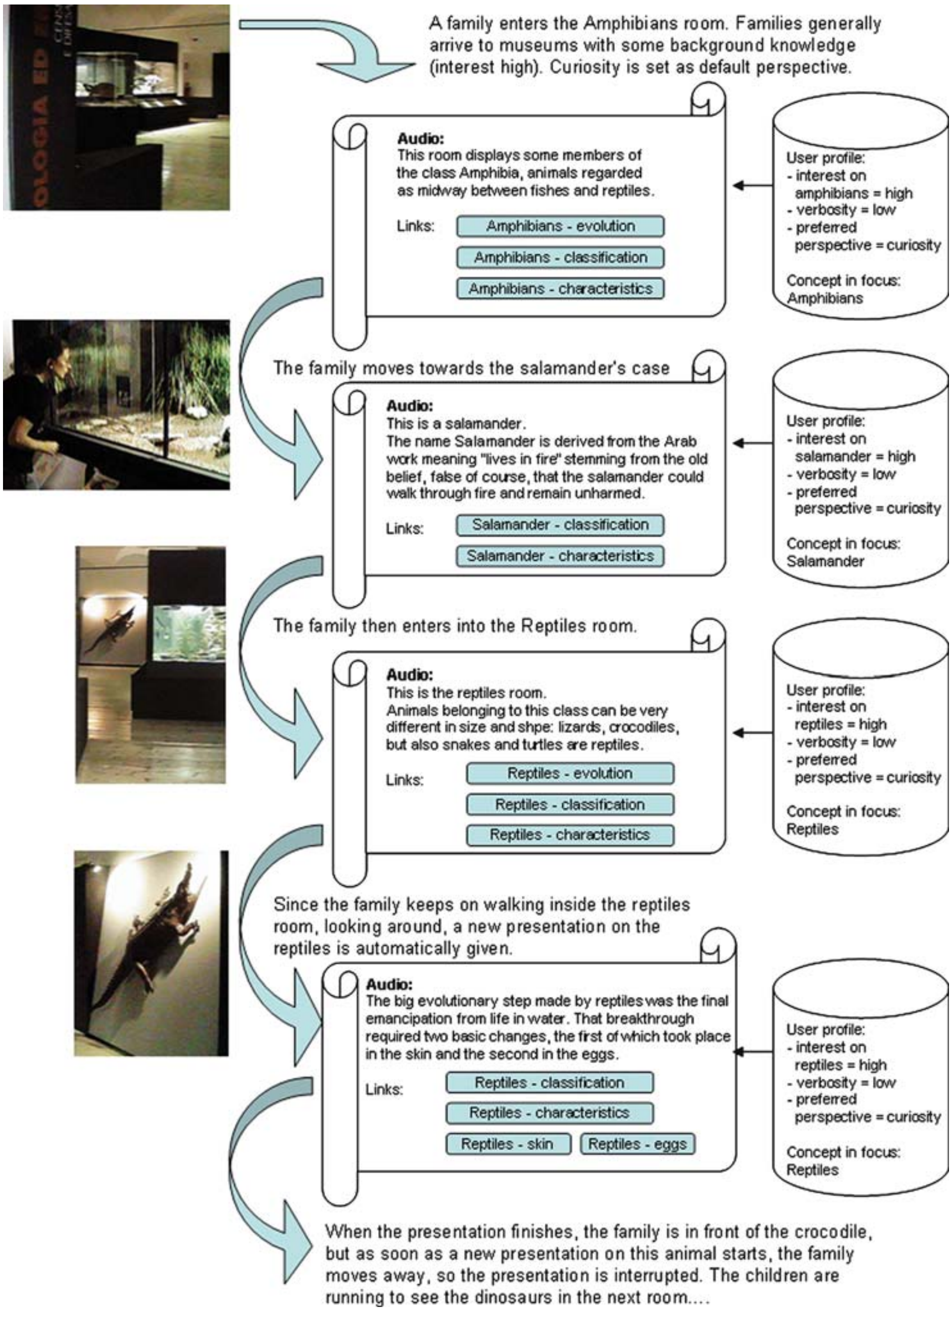

Figure 8. Scenario 3 for user interaction with Hyperaudio.

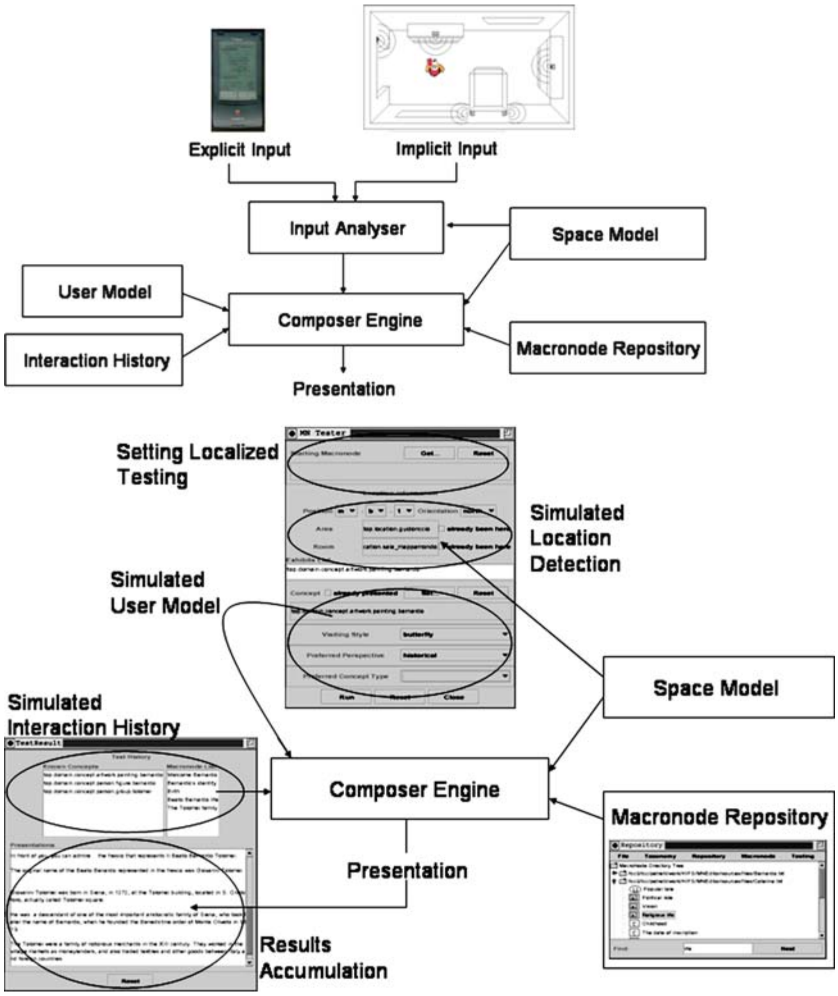

Figure 9. The simplified system architecture (above) compared to the development environment (below).

- 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, no matter whether they are already fully developed or still (hand-) simulated. (Rewritten) Plug-and-play finished components as well as easy disconnection of modules in need of the developer's attention is very important particularly easy removal of modules that need the developer's attention, and easy insertion of completed components, are very important when a development team is involved. When

macronodes and rules were time fine-tuned in the HIPS project, 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 initialization of central parameters. This is particularly important if parameters of simulated modules are likely to change very often, as was the case for the location and orientation of the user in HIPS.
- Quick test-revise cycle: since extensive testing is essential, preparing a test should just require a few mouse clicks. We found it useful to have a graphical panel that allowed us to initialize the context conditions of a hypothetical interaction, to run the system, and to collect the produced output.
- Support localized testing: besides manually setting the context values, it should also be possible to manually set the data to be used in the test. In our system, the initial macronode 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 in combination with a certain data configuration.
- Cumulate the results: it is useful to aggregate 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 adaptive system as a whole and to understand what the user would experience when interacting with it.

A development environment as outlined above allows developers to consider and evaluate many context features, and at the same time to focus 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 is, the more valuable becomes the help provided by the environment, as all aspects of the context are related and can influence each other in negative and unintended ways. Consider a visitor listening to a presentation and moving toward an exhibit. Deciding whether the presentation has to be stopped and a new one started may depend on factors other than the visitor's movements. The decision may for example depend on the type of presentation that is currently playing. If it is about a specific object previously in focus it has to be interrupted (or better shortened until the end of the current sentence) as the references to the physical space are no longer valid. If instead the content of the current presentation is generic, then a full delivery is appropriate (and the new presentation is queued). Using our development environment it is easy to test different rules or combination of rules and to 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 content editing. Content creation must be coupled with immediate testing and revision to guarantee a 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 seven objects, two rooms, and five exhibits) proved to be error prone and time consuming. The cost of hand-writing was prohibitive for the larger HIPS project, and an editing support was deemed necessary.

The need for some editing support when authoring content for adaptive systems has been acknowledged in the literature and recently addressed as an integral part of the development of adaptive hypermedia (De Bra et al., 2003; Weber et al., 2001). While the usefulness of graphical support has been recognized as the complexity increases, the verification of content has only been considered at the level of graph consistency and rule propagation (Calvi and Cristea, 2002). Unfortunately this is not enough when linguistic adaptation is involved: checking for graph correctness would not determine whether a deictic reference was properly applied or the narration was fluent. A human has to systematically check the data structure and its use by the adaptation process. The development environment described in Section 5 can be used to support an author in correctly creating content, by properly supporting the editing of the annotated network, and by 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 (Androutsopoulos et al., 2001): 69 macronodes were created to cover eight exhibits.

Using our development environment has improved the efficiency and effectiveness of adaptive hypertext editing. In the following, we will describe a few lessons we have learned on how to make this tool usable by adaptive hypertext authors outside of the development team (as was the case in the HIPS project). Figure 10 shows the elements discussed below and visualizes 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 regarding the correct compilation. The author can then concentrate on filling them with content and verifying them. This feature is particularly important when the responsibility of creating the data is on authors who are not domain experts, i.e. museum curators. Through templates, the developers can pass the basic knowledge on how the content should 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. In HIPS, 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), a search facility, and user-defined data files turned out to be very handy features for a network composed of hundreds

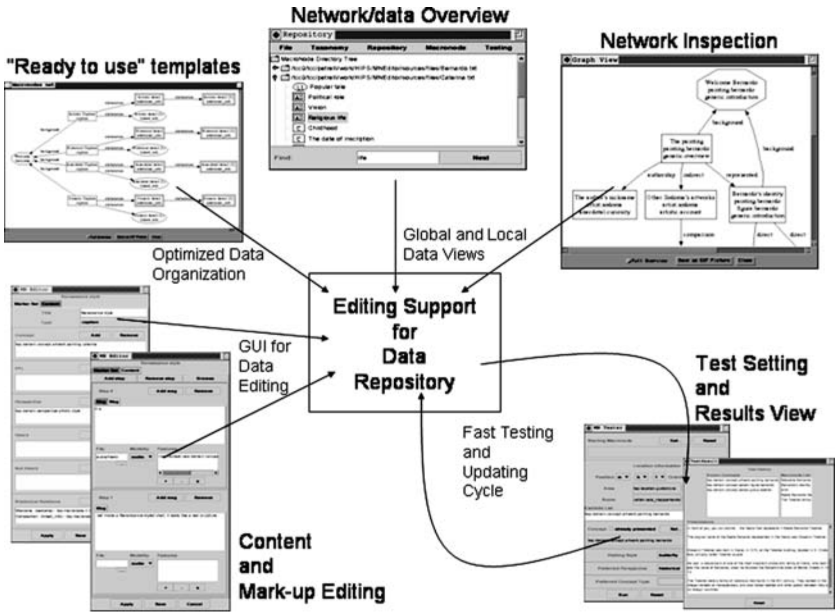

Figure 10. The editing supports components as related to the testing.

of nodes. The possibility of seeing at a glance an idea on, for example, the length of a presentation (i.e. the 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 a chance to be selected and listened to. A further improvement of the graphical interface is the progressive highlighting of the nodes used. This allows the author to quickly check that all nodes can be reached.

- Quick test-revise cycle. Extensive testing during the content editing process is very important. Simple test runs, quick problem identification, and immediate problem fixing have to be supported. The features discussed in the previous section, namely a graphical panel for setting the context conditions, selecting the initial node, launching the system, and collecting the produced output, proved to be a valuable support for fast testing and problem identification from a content verification perspective as well. To support immediate fixing, the panels for testing must be integrated with the editor so that the author can easily switch between the two.

As discussed above, a comprehensive off-line/in-lab testing is essential to assure that the adaptive system is robust enough to go into the hands of users 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 and Paramitys, 2004). However, as the number of mobile adaptive systems increases, there is the need for augmenting the context to location and device (as discussed by Gupta and Grover in Weibelzahl and Paramitys, 2004). New user evaluation methodologies are then needed to test adaptive and mobile guides, particularly when the use is not for work but for leisure (Marti, 2000; Marti and Lanzi, 2001). A first step in this direction can be found in Hatala and Wakkary (2005). The authors 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 for, e.g., the user model. The second phase, the "verification", must be done with real users and should consider dimensions such as variability, sustainability, and evolvability.

7. Related Work on Adaptive Mobile Systems

Adaptive mobile information systems for tourists and travellers is a popular research topic.[9] It seems so promising that a generic user modelling system for tourist applications has been recently proposed (Fink and Kobsa, 2002). This tool is used by the WebGuide system to provide tourists in Heidelberg (Germany) with personalized tour recommendations 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 and Kobsa, 2002).

INTRIGUE (Ardissono et al., 2002) is another tour scheduler. It helps visitors plan tours in Torino (Italy) and its surrounding's and adapts to the needs of a group of people travelling together, e.g. parents and children. Users have to complete an initial "registration form" that provides the system with information on day and time of the visit, and categories and geographical areas of interest. On the basis of this data, INTRIGUE schedules a tour that takes the transfer time into account. It also adapts its layout to the display device, namely desktop PC or WAP phone.

CRUMPET (Poslad et al., 2001) uses a handheld computer (namely an iPAQ) to provide personalized recommendations of services and attractions to city tourists, including 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 a means for fast adaptation but it is not clear. The authors allude to empirical studies for identifying the typical interest profile, but it is not clear whether they have ever been conducted. The interface layout is simple and has been designed with computer/web literate users in mind.

[9] See also Baus et al. (2004) for a selected critical comparison of map-based mobile guides. Note 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 basis of the user's current location, but no adaptation is applied.

GUIDE (Cheverst et al., 2000) implements a location-aware adaptive guide for tourists visiting the city of Lancaster (UK). It recommends sites near the current user position that are open and compatible with the user profile. It can also plan a city tour that spans all sites selected by the user. The order of visit depends on their opening hours as well as the distance and the availability of scenic routes between them. The user interface resembles a Web browser and new information is provided only after a user action. Users are required to fill in a form that asks 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 and Specht, 1999) is one of the preliminary prototypes developed inside the HIPS project. In contrast to 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, and adapts to user interests and knowledge that is inferred from their interaction. Hippie has a browser-based interface that signals the availability of new information by displaying a small blinking icon and playing a "earcon". After 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 for the user's interest profile is available but not mandatory.

While the above applications have different domains, they share the idea of active users interacting with an adaptive guide in 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 PDAs, the interface design is particularly critical and new interaction paradigms need to be explored, as noted by Cheverst et al. (2002). HyperAudio attempted to overcome the limits of screen size 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 performed 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 in HIPS, where a more sophisticated architecture for very fine linguistic adaptation was tried out. A better adaptation to the space and the narration was possible because of a new space model[10] and a deeper discourse context. The space model was finer grained (Bianchi and Zancanaro, 1999), thus allowing for deictic reference to nearby or distant 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 and Zancanaro, 2000, 2001). A new way of modeling users solely on the basis of their

[10]A fine-grained and robust space model is essential to build 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. 2005) opens up a spectrum range of interesting new possibilities.

movements was used in HIPS to adapt presentation length (Marti et al., 2001). Content selection used full models of user interest and knowledge (Opperman and Specht, 2000). Finally, a new graphical interface was implemented to assist users in locating artworks in a room by highlighting them on a 3D user-centred ego perspective[11] reproduction of the room was implemented on the PDA (Gabrielli et al., 1999).

The concept of the disappearing computer has been extended by Zancanaro et al. (2003) who 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 through camera movements and shot transitions using cinematic techniques driven by the underlying content and rhetorical structure of the audio message (Callaway et al., 2005 in press; Rocchi and Zancanaro, 2003; Zancanaro et al., 2003). The video-clips enhance the presentation, help visitors locate the described details in a large and complex fresco, and demonstrate how computer technology can empower and enrich everyday activities. They implement the vision of augmented environments (Wellner et al., 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) on which additional information about the object in view is displayed. A visual augmentation of the museum space is thereby created. 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 of film on the exhibition represents the content. The video clip to be delivered is selected based on the user model, the appropriate order of delivery, and the length of the video. The selected clip is displayed on the private eye with textual and pictorial details.

While Sparacino's system focuses on the visual aspect, the LISTEN project (Zimmermann et al., 2003) 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 modeling 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 tested to integrate the modelling of the user and the modelling of the context (Zimmerman et al., 2005).

[11]The orientation of the room model on the PDA corresponds directly to the user's perspective in the real room.

8. Conclusions

The systems discussed in Section 7 show how adaptive hypermedia are branching out from the narrow path of adapting content and links for users sitting in front of a screen towards a broader adaptation to the interaction context of users immersed in an augmented environment. 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 be increasingly important for successful designs particularly when mobile and ubiquitous computing meets adaptivity. HyperAudio has been one of the few adaptive projects where a user-centred design approach was employed, and is likely to have been the first such project in the area of adaptive and ubiquitous guides. From that experience we learned how a deep understanding of users and uses is essential when designing adaptive systems to be used in highly-constrained conditions, such as set by the performance requirements of a PDA. In this context each design choice has to be evaluated and motivated. In this paper, we have demonstrated that sophisticated and advanced techniques could be inadequate when compared with real use, and how a simpler solution can be equally effective. From our experience, an effective design is based on a few assumptions derived from actual user needs that cover all the aspects of the adaptive system, i.e. the user model, adaptive rules, and annotated data. The creativity of the designers 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. A suitable development environment is mandatory for an iterative design and for the 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 content.

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 envisaging, designing and developing HyperAudio. We feel also indebted to the team of the European HIPS project for the many fruitful discussions and group work that helped our ideas and technologies evolve. The HIPS 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).

Special thanks are due to the Civic Museum of Rovereto and to its director, Dr. Franco Finotti, who inspired the project by showing us the museum and its hardware.

Finally we are grateful to Nigel Ford and Andrew Cox for their encouraging comments on the paper, and to the reviewers and journal editors for the many fruitful and thorough suggestions.

References

Abowd, G., Atckeson, C., Hong, J., Long, S., Kooper, R. and Pinkerton, M.: 1997, Cyberguide: A Mobile Context-aware Tour Guide. *Wireless Network*, **3**, 421–433.

Androutsopoulos, I., Calder, J., Not, E., Pianesi, F. and Roussou, M.: 2001, Multilingual personalised information objects. In: *Proceedings of the International Workshop on Information Presentation and Multimodal Dialogue*, Verona, Italy, pp. 25–29.

Ardissono, L., Goy, A., Petrone, G., Segnan, M. and Torasso, P.: 2002, Ubiquitous user assistance in a tourist information server. 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. 14–23.

Baus, J., Cheverst, K. and Kray, C.: 2004, A Survey of Map-based Mobile Guides. In: L. Meng, A Zipf and 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., Bianchi, A., Marti, P., Not, E. and Sennati, D.: 1999, HIPS: Hyper-interaction within physical space. In: *Proceedings of IEEE Multimedia Systems 99, International Conference on Multimedia Computing and Systems*, 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 Zancanaro, M.: 1999, Tracking users' movements in an artistic physical space. In: M. Caenepeel, D. Benyon and D. Smoth (eds.): *Community of the Future*. *Proceedings i^3 Annual Conference*, pp. 103–106, http://www.i3net.org/ser_pub/publications/proceedings/i3ac99-proceedings.pdf (accessed 15.10.2004).

Broadbent, J. and Marti, P.: 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., Marti, P., Not, E., Oppermann, R. and Rahlff, O.: 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 Awarness, and Adaptivity*, Springer: LNCS 2266, pp. 201–214.

Borun, M. and Dritsas, J.: 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., Karagiannidis, C. and Sampson, D.: 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.

Calvi, L. and Cristea, A.: 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, Springer, Malaga, Spain, LNCS 2347, pp. 79–89.

Callaway C., Not, E., Novello, A., Rocchi, C., Stock O., and Zancanaro, M.: 2005, Automatic Cinematography and Multilingual NLG for Generating Video Documentaries. *Artificial Intelligence* **165**, 57–89.

Carmichael, D. J., Kay J. and Kummerfeld, R. J.: 2005, Consistent modelling of users, devices and sensors in a ubiquitous computing environment. In this issue.

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

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

Cristea A. and Aroyo, L.: 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 Calvi, L.: 1998, AHA! An open Adaptive Hypermedia Architecture. *The New Review of Hypermedia and Multimedia* **4**, 115–139.

De Bra, P., Aqerts, A., Berden, B., De Lange, B., Rousseau, B., Santic, T., Smits, D. and Stash, N.: 2003. AHA! the adaptive hypermedia architecture. In: *Proceedings of International Conference on Hypertext and Hypermedi*, Nottingham, UK, pp. 81–84.

Dieterich, H., Malinowsky, U., Kuhme T. and Schneider-Hufschmidt, M.: 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 Dierking, L.: 1992, *The Museum Experience*. Ann Arbor, MI, Whalesback Books.

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

Gabrielli, F., Marti, P., and Petroni, L.: 1999, Environments as Interface. In: M. Caenepeel, D. Benyon and D. Smoth (eds.): *Community of the Future. Proceedings of the i3 Annual Conference*, Siena, Italy, pp. 44–47.

Hackos, J. T. and Redish, J. C.: 1998, *User and Task Analysis for Interface Design*. New York, 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 Hasan, R.: 1985, *Language, Context and Text: Aspects of Language in a Social-Semiotic Perspective.* Deaking University Press.

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

Houde, S. and Hill, C.: 1997, What do prototypes prototype? In: M. Helander, T. K. Landauer and P. Prabhu (eds): *Handbook of Human-Computer Interaction*. Elsevier Science, Second edition, pp. 367–381.

Jameson, A. and Schwarzkopf, E.: 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 Sampson, D.: 2000, Layered evaluation of adaptive applications and services. In: P. Brusilovsky, O. Stock and C. Strapparava (eds.): 1*st International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, AH 2000, Trento, Italy, Springer: LNCS 1892, pp. 343–346.

Kobsa, A., Koenemann, J. and Pohl, W.: 2001, Personalised hypermedia presentation techniques for improving Online customer relationships. *The Knowledge Engineering Review* **16**(2), 111–155.

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

Mann, W. C. and Thompson, S.: 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 (available from marti@unisi.it).

Marti, P., Gabrielli F. and Pucci, F.: 2001, Situated interaction in Art. *Personal and Ubiquitous Computing* **5**, 71–74.

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

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

Milosavljevic, M., Brusilovsky, P., Moore, J. D., Oberlander, J. and Stock, O. (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*. Academic Press, New York.

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, Cambridge, MA.

Not, E., Petrelli, D., Stock, O., Strapparava C. and Zancanaro, M.: 1997a, Augmented space: bringing the physical dimention into play. In: *Proceedings of Flexible Hypertext Workshop* held in conjunction with Hypertext'97, Southampton, UK.

Not, E., Petrelli, D., Stock, O., Strapparava C. and Zancanaro, M.: 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., Petrelli, D., Sarini, M., Stock, O., Strapparava, C., and Zancanaro, M.: 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 Zancanaro, M.: 1998, Content adaptation for audio-based hypertexts in physical environments. In: *Proceedings of the 2nd Workshop on Adaptive Hypertext and Hyperm*edia, held in conjunction with Hypertext'98, http://wwwis.win.tue.nl/ah98/Proceedings.html (accessed 4.11.2004).

Not, E. and Zancanaro, M.: 2000, The macroNode approach: mediating between adaptive and dynamic hypermedia. In: *Proceedings of the 1stInternational Conference on Adaptive Hypermedia and Adaptive Web-Based Systems*, Trento, Italy, pp. 167–178.

Not, E. and Zancanaro, M.: 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., ODonnell, M., Knott A. and Mellish, C.: 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 Specht, M.: 1999, A nomadic information system for adaptive exhibition guidance. *Archives and Museum Informatics* **13**, 127–138.

Oppermann, R. and Specht, M.: 2000, A context-sensitive nomadic information system as an exhibition guide. In: *Proceedings of the 2ndInternational Symposium on Handheld and Ubiquitous Computing*, HUC 2000, pp. 127–142.

Petrelli, D., De Angeli A. and Convertino, G.: 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., Not, E., Zancanaro, M., Strapparava C. and Stock, O.: 2001, Modelling and adapting to context. *Personal and Ubiquitous Computing* **5**, 20–24.

Petrelli, D., Baggio D. and Pezzulo, G.: 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, Springer, Trento, Italy, LNCS 1892, pp. 202–213.

Preece, J, Rogers Y. and Sharp, H.: 2002, *Interaction Design: Beyond Human-Computer Interaction.* Wiley, New York.

Poslad, S., Laamanen, H., Malaka, R., Nick, A., Buckle P. and Zipf, A.: 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*. Springer, Berlin, pp. 35–51.

Rocchi C. and Zancanaro, M.: 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 Strapparava, C.: 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.is.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., Kuhl H-C. and Weibelzahl, S.: 2001, Developing adaptive internet based courses with the authoring system netCoach. In: S. Reich, M. Tzagarakis and P. De Bra (eds.): *Hypermedia: Openess, Structural Awerness, and Adaptivity*. Springer, pp. 226–238.

Weibelzahl, S. and Paramitys, A.: 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).

Wellner P., Mackay, W. and Gold, R.: 1993, Communication of the ACM, special issue on computer augmented enviroments: Back to the real world **36**(7), 24–97.

Zancanaro, M., Stock, O. and Alfaro, I.: 2003, Using cinematic techniques in a multimedia museum guide. *Proceedings of Museum and the Web*, Charlotte, USA, pp. 209–216.

Zimmermann A., Lorenz, A. and Specht, M.: 2003, User modeling in adaptive audio-augmented museum environments. In: *Proceedings of the 9th International Conference on User Modeling*, UM2003, Springer, Bevin, LNAI 2702, pp. 403–407.

Zimmermann, A., Specht, M. and Lorenz, A.: 2005, Personalization and context-management. In this issue.