1 Adaptive Multimedia Guide

C. Rocchi, I. Graziola, D. Goren-Bar, O. Stock and M. Zancanaro

1.1 Introduction and Motivation

An important goal of museum practitioners today is to find new services that keep the museum a stimulating and up-to-date space for all visitors regardless of age (Verdaasdonk et al. 1996; Laws 1998). From this viewpoint, the information provided to museum visitors is a key component in improving their involvement and thus their attendance (Ravelli 1996). Therefore, technologies are considered a valuable resource for catching visitor attention and stimulating their involvement (Lin et al. 2005). Engineering and HCI researchers have seconded this need with several proposals consistent with this kind of environment, developing museum guides, information kiosks, web services, and progressively more intense multimedia (González-Castaño et al. 2005; Scarlatos et al. 1999). These solutions arose through considerable efforts in studying the museum environment, its communication and interaction potentialities, and interconnections with technology (Ambach et al. 1995). These studies focused on what museum audiences would like to see and learn, which visitor behaviors could be supported by technology, how to offer these technologies to visitors, and visitor attitudes towards these technologies (Silberberg 1995; Crowley et al. 2001). They demonstrated non-homogenous behaviors and interests, which provided an opening for adaptive strategies. Adaptive technologies enable visitors to get personalized information, that is, information related to their particular interest, context and path, without time pressure or the involvement of human guides. They considerably enrich what museums can offer visitors.

In this chapter, we present an adaptive mobile guide, which runs on a PDA and accompanies visitors in Torre Aquila. Communication between the visitor and the system is based on an affective interaction paradigm that allows the visitor to naturally express his or her interest in presentations regarding an artwork. The feedback is exploited by the system to adaptively select the information, according to an inferred user model built up during the visit. We investigated an interaction paradigm that allows the user to express affective attitude towards the presentations proposed by the system. We wanted to verify whether such interaction improved the efficacy of the interface, especially when technology should not hinder the "real" experience, as with museums.

The preparation of adaptive content is a difficult task for multimedia authors, even professionals. For example, writing texts for adaptive systems requires authors to exploit nonlinear writing techniques. In our case, it is even more complicated because the multimedia guide also requires experience in the mixing of text and pictures, along the lines of cinema. Moreover, authors should foster visitor engagement with artworks when organizing the exhibition, identifying communicative strategies, in particular. It is not surprising that many such professionals have a background in humanities (art, literature, and so on) and that they usually do not have knowledge of adaptive systems, and the steps necessary for preparing content for the purpose of personalized composition. In order to manage the adaptive system—and in particular to write and classify contents—the author needs to first understand the vocabulary, the concepts and definitions used by engineers to structure the adaptive system. In the last part of the Personal Experience with Active Cultural Heritage (PEACH) project, we studied this issue seriously, considering how to help authors accomplish the task of authoring adaptive multimedia content. We started from our experience in preparing content for the PEACH mobile guide, and defined a methodology to support authors in preparing the content for adaptive mobile guides. Our first goal was to identify a clear methodology to support authors in preparing "adaptivity-ready" content for a mobile guide. On this basis, we also proposed relevant functionalities of an authoring tool to help authors in writing and managing adaptive content. One of the first papers about adaptive hypermedia authoring is by Hongjing et al. (1998), which introduces a reference model for adaptive applications and which also encompasses the authoring phase. Petrelli et al. (2000) present a graphical interface for preparing content for the framework of the HIPS project. Hyper-Interaction within Physical Space (HIPS, Benelli et al. 1999) was a European funded project aimed at creating personalized presentations of museum exhibits. Recently, a series of workshops about authoring in adaptive applications was organized in conjunction with adaptive hypermedia and education-related conferences.[1]

[1] See http://www.win.tue.nl/~acristea/A3H/.

This chapter is structured as follows: first, we introduce the state of the art in the field of the mobile guides, illustrating the main achievements of the last 15 years. We then describe the evolution of the system—graphical interface and backend—across a set of three prototype-evaluation cycles. For each step, we illustrate the insights that drove the design of the next prototype. Finally, we provide a methodology to help museum curators organize the content of adaptive systems.

1.2 State of the Art

About 15 years ago, Weiser (1991) proposed a set of ideas about the future of computers in the next century. “Personalization” was one of the key words mentioned in that paper and later was identified as a key feature of ubiquitous computing systems (Abowd et al. 2000). In a series of attempts to port “office applications” to mobile devices, researchers realized that new features could be explored, including information sharing (Aoki et al. 2002), location and context awareness (Want et al. 1995; Cheverst et al. 1999), and adaptivity (Rocchi et al. 2004).

From the second half of the 1990s on, there have been many efforts to explore the potential of mobile systems. CyberGuide was a large project which aimed at exploring the use of a mobile, hand-held, context-aware tour guide (Long et al. 1996; Abowd et al. 1997). The authors illustrate both indoor and outdoor prototypes, explaining the issues of detecting user position and orientation, and discussing the choice of hardware, the appropriateness of communication media (audio or video), and the methodology of map representation.

The HIPS project focused on hyper-interaction (Benelli et al. 1999), which evolved from the earlier HyperAudio experience (Not et al. 1998; Petrelli and Not 2005). A novel aspect, with respect to previous projects, was the overlapping of contextual and personalized information on top of the physical space. The user experience is augmented because spatial and informational navigation occur at the same time. The system, set up in Museo Civico (Siena, Italy), used a hand-held device to support visitors in moving around, and in seeking information and guidance. Personalization was based on user position, and on interaction with the PDA and the surrounding physical space.

GUIDE is another successful mobile system (Cheverst et al. 2000). It supports tourists visiting the city of Lancaster (UK). Combining mobile technologies and wireless infrastructures, it tailors information to the user’s personal and contextual situation. Its design, carried out in collaboration with experts in the field of tourism, is particularly valuable and the insights gained during evaluations brought interesting changes to the prototypes.

Survey papers have helped assess the development of research from different perspectives. For instance, Kray and Baus (2003) present a survey of mobile guides, both prototypes or commercial, whereas Raptis et al. (2005) attempt to classify current practices in the design of mobile guides for museums.

Some systems featuring mobile guides are DiscoveryPoint, the Genoa Aquarium Guide, and SottoVoce. The first is a remote control-like device that allows users to listen to short stories related to an artwork; it is installed at the Carnegie Museum of Art in Pittsburgh (Berkovich et al. 2003). It is an audio system with a special speaker which delivers pinpointed audio and can be heard near the work of art.

Another PDA application has been tested at Genoa's Costa Aquarium (Bellotti et al. 2002). The basic elements of the interface are multimedia cards, each corresponding to a presentation subject such as a particular fish or a fish tank containing several species. Each multimedia card provides users with content. Touch-screen buttons allow control of content presentation and navigation between tanks.

Grinter et al. (2002) report on an interesting study about the SottoVoce system which was designed to promote a shared experience during a visit to a historic house. The system supports shared playing of audio content between pairs of visitors, each using a PDA. The paper reports interesting findings on how this technology helps to shape the experience in the museum (shared versus individual use of the device).

Multimedia guide literature also includes research concerning the development of architectures for context-aware applications. (Dey et al. 2001) present an interesting attempt to define the notion of context and introduce a conceptual model with a set of methods to help drive the design of context-aware applications. Their proposal is a computational framework to quicken the implementation and prototyping of context-aware applications.

Efstratiou et al. (2003) introduce a new platform to support the coordination of multiple adaptive applications. Coordination can be specified in terms of policies (they also present a formal language to define them) which allow adaptive (re)actions on a system-wide level to be described. For instance, it is possible to define the level of system intrusiveness, for example, whether to notify the user of system actions.

Seamless connection between mobile devices is a recent research issue. Krüger and colleagues focused on providing user-adapted seamless services in different situations. They worked on a route planning scenario where a user is supported by a central system while using three devices: a desktop route planner, an in-car navigation system, and a pedestrian navigation service running on a PDA (Krüger et al. 2004).[2]

Evaluation studies provide insights on both design and reimplementation of prototypes. In addition to previously cited works, the added value of systematic evaluation is shown in Bohnenberger et al. (2005), which reports on improvements brought to a mobile shopping guide after an iterative evaluation cycle. In particular, the authors focus on usability issues of the PDA interface with respect to the task to be accomplished (buy items in the minimum possible time) and on the accuracy of the system in supporting the user.

1.3 History of the System

At the beginning of the project, we stated the requirements of a mobile guide. Unlike many mobile devices (e.g., cell phones) our guide had to be *wake-up-and-use* since it is not meant to be used daily. This means it had to be intuitive and clear, with almost no need for explanation. Second, it had to be nonintrusive, that is, it could not interfere with a visitor's enjoyment of artwork. Third, we wanted to experiment with a new communication paradigm based on *delegation* and *affect.* In many "information-seeking" scenarios people typically look for information (think of Web queries with search engines). We consider the museum experience as something that has more to do with entertainment than with information seeking. The system's proactiveness in delivering an appropriate presentation at the right time is essential. Therefore, we devised interaction based on delegation: visitors do not ask for information about an artwork, they rather signal liking or disliking (that is affect towards the artwork and the information presented by the system) and, almost implicitly, request information from the system. This feedback from the visitor can be seen as a sort of nonverbal backchannel gesture that the system takes into account in selecting information for that specific user. This introduces the fourth requirement: the system had to be adaptive, that is, it had to provide personalized information according to the user profile built during the visit. We also wanted to experiment with affective interaction, believing it might improve system usability and acceptability.

In the design of the system, we kept in mind four fundamental adaptive dimensions which were also used for the evaluation phase:

[2] See Chap. 7 on seamless presentations across devices.

- Location awareness: the system initiates interaction when it detects the visitor at a given position. This feature is enabled by means of sensors located close to each work of art.
- Follow-up: the system arranges the following presentations on the fly according to feedback received from the visitor through the interface.
- Interest: the system selects content appropriate to (estimated) user interests.
- History: the system appropriately refers to previously seen items or points to artworks closely related to the one currently visited.

From the design requirements, we moved along iteratively, through quick prototyping and small evaluations. This drove the evolution of the system—both the graphical interface and the backend—to the final prototype, which is based on a *like-o-meter* widget and a shallow semantic network for dynamically modelling visitor interests. In this section, we shed more light on the steps in this evolution.

1.3.1 First Design

The first design is based on a two-button interface, as illustrated in Fig. 1.

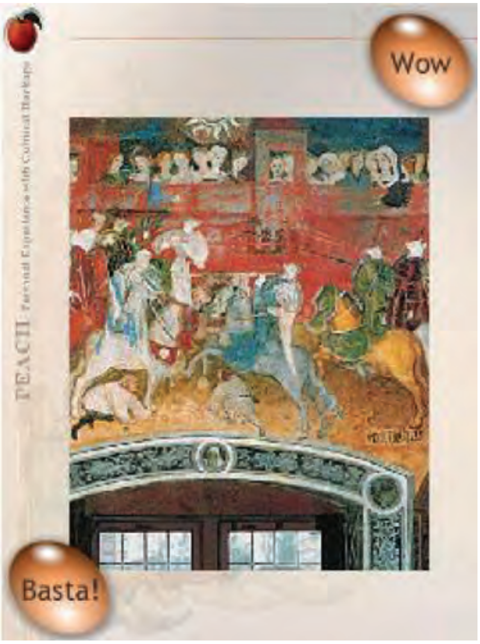

Fig. 1. The first prototype of the interface

Pressing the WOW button is meant to express a visitor's interest whenever she is impressed by a fresco or by a specific topic related to it. The BASTA! button, on the other hand, was to be used for lack of interest in current topic. As a side effect, the BASTA! button stops the current presentation. It is worth noting that a presentation can also be stopped by moving away from the current fresco to approach another one.[3] The central part of the screen was used to show a video presentation.

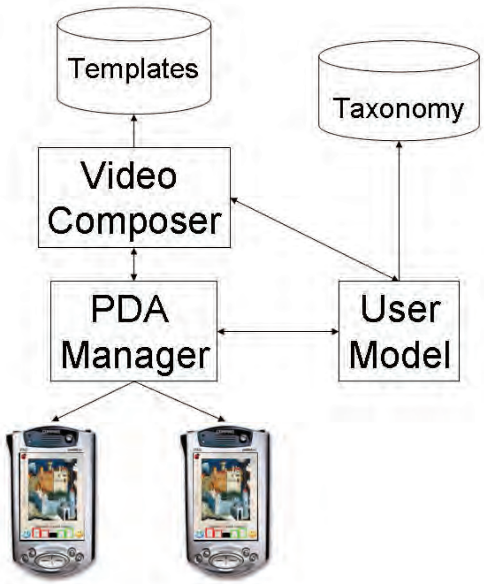

Fig. 2. Architecture underlying the first prototype

The underlying architecture that supports the first prototype consists of the following components (Fig. 2):

1. A video composer (VC) for dynamically composing presentations from a repository of templates according to the user interests stored in the user model.
2. A user interface manager (UI) for catching user location, interaction through buttons, and dispatching the presentations generated.
3. A user model for receiving the messages from the UI, for computing the current interest level of the visitor on each of the topics, and for propagating interest level according to a topic's taxonomy.

[3] This feature is "discovered" by the visitor during the visit.

4. A tree-based topic taxonomy which represents the contents of the pictures at three levels of abstraction (from specific content to abstract concepts such as “aristocratic leisure activities”).

Templates, encoded in the XASCRIPT formalism, allow definition of a set of possible video documentaries from which user-tailored presentations are dynamically assembled (Rocchi and Zancanaro 2003). A template contains a set of instructions for image selection (or audio clips), and variable settings. By means of a merging mechanism, the system uses a single template to compose different video documentaries, conditioned by the user profile built by the user model (see Chap. 2).

1.3.2 Initial User Studies

Using the prototype described above, the first user study was set up in the real scenario, Torre Aquila, and included eight visitors. In the tower, we installed:

- A laptop, running the server side of the system.
- An access point, to allow wireless communication between devices.
- A PDA, running the interface shown in Fig. 1.
- Four infrared beacons, one for each fresco, to detect user position.

Users were “real” visitors, recruited at the entrance. They were given a short verbal introduction about the guide, followed by a real visit of four frescos (of eleven) exhibited. The experimenter observed the users during the visit. At the end of the visit, an informal interview was performed in order to assess the perception of the four adaptive dimensions: location, follow-up, interest, and history.

At the end of this evaluation cycle, both looking at interviews and considering experimenter observations, it was apparent that users did not clearly understand the graphical interface. In other words, during the experiments to investigate the effects of adaptivity on visitors, we found that the system was not usable. The WOW button was the source of many misunderstandings of the whole system. Sometimes it was pressed to initiate the interaction, although this is not needed at all.

To study usability and user perception of adaptivity in more depth, we resorted to an “action-protocol and retrospective-interview” qualitative study, targeting the expression of the affect and the delegation-of-control paradigm. The main difference of this methodology with respect to think-aloud is that the user does not provide her comments during the execution of the task, but later on, while she and the experimenter are watching a video recording of the interaction (Van Someren 1994).

The study was conducted on three subjects, in a room equipped with posters of the originals frescos, and with sensors to detect the positions of the subject with respect to the frescos. Although small, this is deemed to be a reasonable number of users for an initial investigation according to Nielsen (1993). The room was also equipped with two cameras to record both user behaviour and speech during the visit. The subjects performed a visit by using the guide; the study was limited to four frescos. At the end of the visit, each subject was interviewed by the experimenter while both were watching the video of the visit. The interview focused on the subject's understanding of the WOW button. The interview was recorded, providing important additional material for the research and design teams to discuss during the post-study phase. The results of this study revealed very interesting findings:

1. Loading the information at the beginning of each presentation takes few seconds, which was too long for visitors. In that amount of time, users get disoriented because they do not know what to expect. They do not understand their role with the system, and the WOW button is the only action they can perform during that time, probably in hope of getting information. The instructions did not prepare the users well for that situation.
2. At the end of the presentation, the system stops and shows a default screen shot, delegating the continuation of the visit to the user. This incoherence in the conceptual model confuses the user again, who presses the WOW to get some instructions but, instead, gets another presentation because the interest model was reinforced.
3. When all the content about a fresco was presented (including extended presentations retrieved by the user model) the system expects the user to move on. However, the user does not know the system status, therefore, she still expects something to occur. Again, the only available button is WOW.
4. The BASTA! button is interpreted just as a stop because it causes the system to stop the presentation. This is also incoherent with respect to the design guidelines, because the system should not enable to take an action, but just to express feelings (in this case "don't like").

In the end, the WOW button is often used as the resource of "last resort" to communicate with the system in case of problems. This shows that the intended conceptual model was not clear to the subject, and that the system often has incoherent or unexpected behaviour. In particular, the presentation should not abruptly stop. Rather, the user should be invited to move to another fresco. The system should give feedback about its own status, and inform users about its estimates of user interest; Kay (2001) calls this feature scrutability. It should skip uninteresting presentations, and focus on more interesting ones, or suggest moving to other frescos.

1.3.3 Second Redesign

Based on the results of the pilot study, we redefined the initial requirements for the new prototype:

1. The UI should clearly and intuitively enable the user to express her feelings towards the exhibits during the museum visit.
2. The UI should be coherent and consistent, reflecting the delegation of the control interaction paradigm on all four adaptivity dimensions.
3. The UI should be proactive in order to avoid user disorientation, even if users are told that the system tracks their position. For example, the system should signal when a new position is detected.
4. The UI should give visual feedback to the user, relating its understanding of user interest and its current status (such as preparing presentations for display), without disturbing user attention.
5. The information provided by the system must be structured differently, to allow different degrees of personalization.[4]
6. If the visitor does not express any feeling (that is, she never presses any button), she should receive a reasonable amount of information about the museum's exhibits.
7. Each presentation must have a title for display during the playing.

These requirements led to the prototyping of the second interface, shown in Fig. 3, which features a new widget (at the bottom), called the *like-o-meter*. It substitutes the two previous buttons and aims at better conveying the delegation paradigm. This widget allows the user to express her interest towards the current presentation, by moving the slider towards the smiley face (two degrees of liking), or state "I don't like it", by moving the slider towards the sad face on the left (two degrees of disliking). The presentation title appears close to the widget, helping the visitor to remember what she is "scoring". The feedback from the visitor is taken into account

[4] This requirement came from a strategic research decision and not from pilot study evidence.

by the system, which activates a propagation mechanism at the end of each presentation to update the interest model of the visitor.

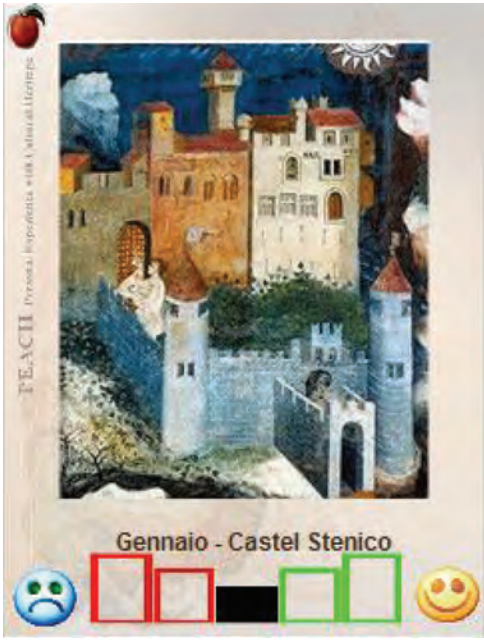

Fig. 3. The second interface prototype, based on the *like-o-meter*

The propagation is performed on a network of templates, as shown in Fig. 4. In fact, in this redesign, we also made modifications to the backend of the system. We decided to drop the taxonomy of topics, in favour of a content-based classification of the templates. Instead of using a user model, reasoning on the structure of the domain at different levels of abstractions, we wanted to experiment with an explicitly linked network of contents, classified on the basis of their communicative functions (see Sect. 1.4).

The template network is organized in nodes, each having an interest value that is zero at the beginning. The interest value assigned to each template ranges from –2, meaning current lack of interest or likely future lack of interest, to +2, meaning strong interest or likely future strong interest. The links between nodes are created by the author network. Each link is a shallow semantic relation which can be paraphrased as “related to”.

Nodes can be of five types:

1. Introduction: contains a general overview of an exhibit
2. Abstract: quickly describes a part of an exhibit
3. Content: extends the presentation of an abstract
4. Follow-up: describes general themes shared by two or more exhibits (e.g., hunting in the middle age)
5. Conclusion: tells the visitor that the presentation of an exhibit is over

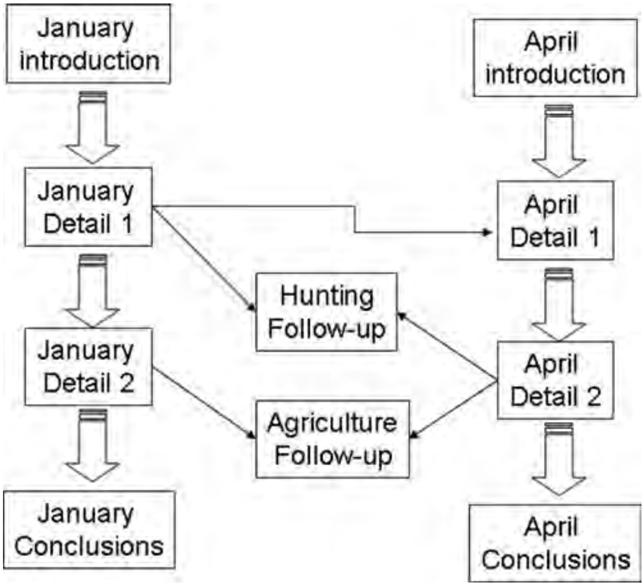

Fig. 4. An example of the template network

As feedback is received through the interface, the system updates the interest value of the current template and also propagates such information to its connected templates, according to the following dependency relations:

1. Introduction affects abstract: a positive or negative degree of interest, expressed when the system presents the general introduction to an exhibit, is propagated to all the abstracts pertaining to the same exhibit. The abstract is the basic information on an exhibit with respect to a certain topic.
2. Abstract affects content: the degree of interest towards the abstract updates the value of its related content (that is, a more detailed

description of the exhibit with respect to the abstract). When a content template has an interest value greater than zero, it is selected and presented to the visitor right after the abstract.

3. Content affects follow-up: the degree of interest towards contents affects its connected follow-ups (if any). Follow-ups are selected and proposed to the visitor when they have an interest value of +2.
4. Follow-up affects content: the degree of interest towards follow-up is propagated to all of the connected contents. Follow-ups act as bridges by propagating the interest values on one exhibit to other similar exhibits.

The selection mechanism is based on an algorithm that ensures that:

- The introduction is always the first presentation.
- Abstracts are always presented and sorted according to the interest value (from the greatest to the least).
- Contents are selected if they have a value greater than zero.
- Follow-ups are selected if they have a value greater than one.
- The conclusion is always the last presentation for an artwork.

The algorithm is iterative and starts from a skeleton, where the introduction is the first element and the conclusion is the last one. All the elements in the middle are adaptively arranged according to the visitor's interaction with the system. The overall goal of this mechanism, coupled with the visual feedback provided through the like-o-meter, is to give the visitor a clear indication that her actions do have an effect on the presentation, while avoiding the interpretation that her own actions are an explicit request for more information. As long as she remains at the same location, the system plays all the presentations, until the conclusion.

Through the widget, the system also informs the user about its own assessment of her interest on the current presentation, by presetting the like-o-meter. Thus, this widget is at the same time an input device and an output device that the system exploits to inform the visitor about the user model. This satisfies the necessity for the user to control the user model, as pointed out in Kay (2001).

The system performs another type of content adaptation with respect to the history of the interaction. The UI component keeps track of the user visit, enabling the comparison of the current presentation with previously seen ones. This is enabled by the merging mechanism, which can dynamically select presentations that explicitly refer to the visitor's history (such as "as you have seen before…").

Before making all these modifications, we created a partial mockup using Macromedia Flash to get an early evaluation of the new concepts and architecture. The mockup employs a hand-coded template network and a hand-coded propagation mechanism for one specific fresco only.

In this third study, we focused on how well the users are able to recognize and use the like-o-meter widget. Using the following questions, we investigated whether the like-o-meter properly communicates its meaning:

- Do the users recognize it as a scale with two negative and two positive values plus a neutral position?
- Do they understand that the position of the place card on the bar signals to the system their level of like or dislike of the current presentation?
- Do they understand that their expression of liking/disliking relates to the current presentation and not the whole exhibit or to a specific utterance?
- Do they notice the user model feedback on the like-o-meter bar, recognizing it as a consequence of their previous expressions of liking/disliking?
- Do they recognize that when the system provides more information this occurs because of their expression of interest?

The user study consisted again in an action protocol with retrospection. It was conducted on two users acting in the same room with the panels reproducing the original fresco exploited in the previous user study. The experimenter first presented an introduction of the museum setting and showed a copy of the fresco used in the experiment. Then, she quickly demonstrated the functioning of the system.[5] We decided to employ a task-based scenario in order to focus on the issues relevant to understanding of the metaphor. Another reason for this approach was that the new backend system was completely different from the old one (and not yet completely implemented). The task to the user was to signal interest during the presentation and a slight dislike during the description of the first detail. They had to be enthusiastic about the description of the second detail, and to stop the presentation during the description of the third detail. This task allowed us to check for proper understanding of the like-o-meter while assuring that the interaction could be handled by our partial mockup. The results of this user study were quite encouraging, showing a high degree of understanding and satisfaction by the users. The generalization of the findings to long-term effects of the expression of liking could not be reliably

[5] She simply explained that the like-o-meter is a way to tell the system about preferences. She did not mention any side effect caused by such interaction.

done at that time, given the limited possibility of interaction our mockup allowed.

From the responses and comments of the participants the following considerations emerged:

- The participants were able to communicate their interest to the system, correctly using the like-o-meter.
- The participants recognized that positions +1 and +2 caused more information to be provided in the general visit.
- The relationships between the standard (abstract) and the in-depth presentations (content and follow-up) were clear to the subjects. However, given the limitations of this small study, we cannot reliably conclude that the delegation metaphor was properly understood by the subjects, though this seems likely given the available evidence. In particular, we cannot reliably conclude that they fully realized that their expression of interest on the current exhibit also affected the presentations to come.
- The understanding of the meaning of the moderate disliking (i.e., position –1) is somewhat poorer than that of the liking. Apparently, the users come to expect that the expression of a moderate disliking should cause the system to provide less information. In our current system, on the other hand, the expression of a moderate dislike only changes the user model and does not affect the current presentation.
- The users did not expect that the neutral position of the like-o-meter could be selected, and expected that a single button press would have moved the slider straight from –1 to 1. Actually, we realized that the neutral position may have two distinct meanings: it communicates a degree of liking which is neither positive nor negative, while also corresponding to a lack of information about the user interest. Both of our users seemed to stick to the second meaning, expecting that only the system would be allowed to use the neutral position.
- One participant clearly noticed the feedback of the user model and understood that it was related to her previous behaviour. Both participants understood this feedback as a system initiative.

Summing up, the participants were able to properly carry out the task with a reasonable understanding of the conceptual model of the system. They both agreed that the interface was easy to use and that their expectations about the interest model were fulfilled. This encouraged us to go on with the implementation of the second prototype. In the meantime, we tried to work out the misunderstanding about the neutral position, and we restyled the interface graphics with the help of a professional designer. This led to the last interface, shown in Fig. 5. Apart from stylistic changes, the like-o-meter has been rendered with a needle, buttons have been enlarged to facilitate clicking, and the neutral position has been implemented as usable only by the system, for conveying its assumption about the user's neutral attitude towards the presentation.

Fig. 5. The last prototype of the interface

1.4 Authoring Guidelines

During the final phase of PEACH project, the museum adaptive guide was ready, but most of the content was still missing. So we hired two professional multimedia authors to prepare new content for the guide. Both authors used to work in a famous Italian publishing company. One was an expert documentary maker who also worked for the Italian national TV; the other was a multimedia designer who realized many CD-ROMs related to art. Neither is an engineer nor has a computer science background. The authors needed four months to complete the task (approximately one hour of adaptive multimedia content). During this period the authors needed several sessions of training to acquire basic skills about adaptive components. In particular, they first had to understand the vocabulary, the concepts, and the definitions used by engineers to structure the adaptive system (see the final prototype description above).

The authors were quite skilled in digging out information on artworks from the literature, but they had to structure the content as templates, for example, deciding what contents to associate with each template type. They resolved this issue by deciding that each template had to address a communicative function, as follows:

- Introduction: which exhibit is this? Who's the author?
- Detail: what are the issues or the particulars of the work-art that we would like tell?
- Abstract: what's this part?
- Content: tell more about the previous item.
- Follow-up: give a general context for this information.
- Conclusion: the presentation concerning this artwork.

In order to acquaint the author with the behaviour of the system, we explained it by examples, in a scenario-based approach, illustrating all the possible system "reactions" according to different user interaction. Authors started preparing contents from the "default" sequence, by putting together texts and images of the basic set of information for each exhibit; they would write scripts in two columns, audio and video, like cinema script-writers (Table 1).

Table 1. An excerpt of script

Video track	Audio track and voiceover
	Music, 15 sec. from CD 2, track 3.
Display whole picture of July. Focus on the blue sky and the sun.	The sun stands out of the blue sky. It is July sun, in Leo constellation,
Whole picture of July.	full of activities; almost all the fresco depicts work scenes.
Zoom on three scythe men on the left.	Peasants are busy haymaking
Zoom on four people with hay forks.	and raking up.

Further comments or instruction for synchronization were written in a third column. Once developed and tested the default sequence, authors—with our help—started implementing variations, that is, different *paths* of presentations in order to accomplish the adaptivity. We observed some difficulty for the authors in organizing adaptive variations. At this stage, a graph representation helped us to explain system behaviour in particular situations to them. For example, the graph in Fig. 6 helped us to explain the notion of fork, which can generate two different paths, e.g. *(shot1, shot2, shot4)* or *(shot1, shot3, shot4)*. The condition to choose between shot2 and shot3 relates to the profile built by the system. For instance, shot2 might contain a reference to a previously visited item (e.g., *as you have seen in …*) and shot3 might contain a suggestion (e.g., *as you will see in …*). These are simply textual variations, but authors are also allowed to modify the video part and include new or alternative visuals.

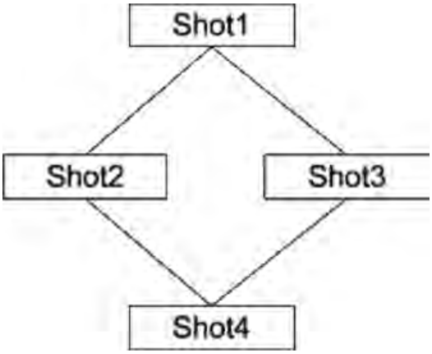

Fig. 6. A graph representing alternative paths

Images were selected from a corpus of high-resolution pictures provided by the museum. Authors were free to crop, resize, and choose the portions of picture they preferred to use in building the visual part of the shots. Texts were synthesized using a professional synthesizer, and Macromedia Flash was used to edit the presentations.[6]

After the preparation of the first exhibit, we tested the content, to find out how effective our training had been. The testing consisted in simulating different interactions with the system and its associated system output. This phase was helpful to identify possible repetitions or inconsistencies in the visual parts. The test was done on a PDA to verify how effective the choice of pictures was and to check the quality of the video on a small screen.

According to the experience described above, we proposed a set of guidelines for preparing the content of an adaptive application. For each phase, we also proposed the functionality to be implemented in the authoring tool:

[6] http://download.macromedia.com/pub/documentation/en/flash/mx2004/flite1_1_authoring_guidelines.pdf.

1. Identify the salient parts of each exhibit.[7]
2. Take notes of the content of each unit. Authors should start from the "mandatory" pieces of information which have to be presented (introduction, abstracts, conclusion). Taking notes helps authors to sketch out the general structure of the content. If notes are organized by keywords, the system might help to highlight connections between exhibits.
3. Find connections between (parts of) artworks. For example, two pictures representing the same subject (or a similar one), can be connected. Each of these connections identified during this phase can potentially be a follow-up.
4. Prepare the scripts. In our experience, organization by columns proved successful and clear (Table 1); indeed, this method is also used by cinema scriptwriters. During this phase, it is also important to test possible combinations of shots to identify potential repetitions or incongruent sequences. "Debugging" came too late in the first prototype implemented.
5. Edit the presentations: assemble selected and elaborated images, and synthesize the audio.
6. Test on the PDA. The final test should be done on the PDA in the real scenario.

1.5 Conclusion

We have illustrated the evolution of an adaptive multimedia mobile guide developed in the framework of the PEACH project. We started by describing the motivations behind the requirements of an ideal adaptive guide. We introduced a new communicative tool—the video documentary—which is adaptively composed to help visitors identify exhibit details mentioned in the audio counterpart and, moreover, to ease finding the relationship between new and already presented information (Chap. 2).

Through an iterative process, we described the results achieved at each stage of the user-centered design. In particular, we focused on the evolution of the interface, and its acceptability and usability after each redesign. We also described the changes made to the system backend. The main result of the process is the final prototype, an interface based on affective interaction, graphically conveyed through the like-o-meter, that allows

[7] In our scenario, we identified the salient parts of each exhibit, that is, the parts that were generally described in the texts provided by the museum curators.

visitors to express their attitude towards the exhibit and explicitly signal their interests to guide the adaptation mechanism of the system.

Finally, we described our experience in training two authors to prepare adaptive content. By considering and analyzing the experience, we devised a set of guidelines aimed at helping the authors prepare and structure content for adaptive mobile guides.

1.6 Acknowledgments

The authors wish to thank Franco Dossena and Adriano Freri for their work on content preparation in relation to the Torre Aquila frescoes.