SECTION II: INTERFACE DESIGN AND PRESENTATION

Person-Oriented Guided Visits in a Physical Museum

E. Not, D. Petrelli, O. Stock, C. Strapparava, M. Zancanaro
IRST, Cognitive and Communication Technology Division, Italy

Abstract

The "ideal guided visit" to a space (as an exhibition, a museum, an open-air exposition distributed in different areas of a city, an archaeological site and so on) allows the visitor to organise the tour throughout the different areas according to their own needs or preferred criteria: for example, looking for paintings coming from the same geographical region rather than paintings produced in the same period (following a geographical thread rather than a chronological one). However, when organising the physical layout of the exposition areas a specific criterion for items placement has to be chosen, according to a "default" perspective of information presentation, either chosen by an architect or imposed by geographical constraints. It may happen that the default physical organisation does not meet directly the visitor's expectations, possibly making it difficult to build a personal route.

We shall discuss the development of a system able to organise the presentation of a museum contents taking into account the visitor's needs and the layout of the physical space. The system guides the visitor through the museum and provides information by means of audio messages: the visitor can get instructions for finding the objects of interest, hear descriptions with references to other interesting objects in the museum, ask for additional information and receive suggestions on next steps. Therefore, the visitor is provided with a personalised guide for exploring the physical and information space.

Keywords

adaptive hypertext; intelligent interfaces; user models; portable digital assistants

1. Introduction

Artificial Intelligence, and in particular intelligent interfaces, may open new scenarios for tourism and fruition of cultural heritage. The main prospect made possible by the adoption of these technologies is to move from the current mass-oriented approach to an approach oriented toward the individual. This has at least two facets: on the one side systems will have to take into account the specificity of the user with his own interests, idiosyncrasies, and so on; on the other the user will be the main agent in his exploration, he will take the initiative and exploit the high level of interactivity that will be available in intelligent interfaces.

With this vision of the future, at IRST we have worked for several years at developing a prototype in which some advanced research concepts were put together in a complex system demonstrable in a realistic context. Before focusing on the specific object of the present article, we would like to briefly review the characteristics of that project, called AlFresco [ST93], [Sto94].

The AlFresco system is a natural language dialogue system for a user interested in Fourteenth Century Italian frescoes and paintings with the aim not only of providing information, but also of promoting other masterpieces that may attract the user. It runs on a workstation connected to a videodisk unit and a touch screen. The particular videodisk in use includes images of frescoes and monuments. The system, besides understanding and using natural language, integrates it with hypermedia both in input and output. The user can interact with the system by typing a sentence, navigating in the underlying hypertext, and using the touch screen. In input, our efforts have been focused on combining the interpretation of linguistic deictic references with pointing to images displayed on a touch screen. In output, images and generated text with buttons offer entry points for further hypertextual exploration. The result is that the user communicates linguistically and by manipulating various entities, images, and text itself. In sum AlFresco yields an environment suited for exploration of a large information space where the user may integrate moments of goal-oriented investigation and browsing.

A system like AlFresco can constitute the basis for exploration at home. But we believe nothing can take the place of a visit to the "real thing". The emotion is different. So the whole cultural experience can be seen as encompassing three phases, each of which can take advantage of technology: support before the visit, during the visit and after the visit.

Here we shall describe an approach to support the visitor *during* his visit. The perspective is that the visitor moves in the physical space while seeking information and guidance. The user is not moving in virtual space but in a real space (such as a museum) augmented with a personalised informational dimension. The two spaces are integrated, producing a new way of exploring the culture heritage. The solutions presented in the paper are currently being developed inside a project called HyperAudio [HA]. Further challenging research issues emerging in this interaction scenario will be investigated inside the HIPS [HIP] European project.

2. THE IDEAL GUIDED VISIT

Usually, a tourist experience begins before reaching the actual cultural/tourist site. In fact, the interest for a particular place or masterpiece often begins with reading a book or an advertisement or talking to a friend; it grows by reading preparatory material; it may continue after the visit by purchasing and reading the catalogue or related books. The knowledge acquired before the visit is used to better enjoy the actual tour and is completed later on by further study.

Information technology can support the visitor's thirst for information in different ways either before, during and after the actual cultural/tourist visit, helping the visitor to get the best out of this experience.

2.1 PREPARING FOR THE VISIT

At present, the most common way of acquiring knowledge on cultural exhibitions or tourist sites still consists in reading books on the subject or advertisement material prepared by the organisers, talking with more informed friends or consulting experts. However, these traditional forms of knowledge acquisition have also technological counterparts that are becoming more and more used: for example, CD-ROM based systems; digital libraries accessed through Internet; intelligent information-providing systems. Within this latter class, many systems have been developed exploiting the more advanced findings of the AI field, especially in the context of adaptive and adaptable systems. In the cultural domain related to museums, for example, we find systems like the before mentioned ALFRESCO, ILEX[1] [KMOO96], [MOOK97] or PEBA-II[2][DM96], just to mention a few.

More generally, the use of technology has a lot of obvious advantages, e.g. world-wide access and adaptive presentation of information [Fle97]. But, there is more than that: the information exploration performed at home by the visitor-to-be can be used to collect information on his previous knowledge and preferences which will be effectively exploited during the visit. For example, to allow on-site systems (inside the real museum or the actual place visited) to be more effective in proposing interesting information.

2.2 CONTINUING EXPLORING INFORMATION AFTER THE VISIT

Usually, visitors bring home books, CD-ROM or video catalogues related to what they have seen. And they often browse through the collected material in search for what they have learnt, what they liked most and further material to enrich the knowledge they have acquired. These types of traditional *static*[3] catalogues usually gather a lot of interesting information related to all the objects displayed in the exhibition. The material is typically organised by domain experts according to a chosen perspective (the leitmotiv that inspires the conceptual and physical organisation of the exposition) and possibly educational strategies.

Recent advances in the field of Artificial Intelligence provide valuable background for the development of new-generation catalogues, complementary to the more traditional ones, summarising the information explored by the visitor according to a meaningful presentational structure and tailoring it to the single visitor's personal traits. This new type of catalogue could be automatically created on the basis of a user-model, built before and during the visit, and completed by interacting with a workstation connected to a multimedia data base. Opposite to the traditional one, we call this type of catalogue *dynamic*.

2.3 THE ACTUAL VISIT: MOVING IN AN AUGMENTED SPACE

In an "ideal guided visit" to a space (as an exhibition, a museum, an archaeological site, a city, and so on) what a visitor would like to have is a flexible companion that helps him in visiting the museum or town as an information space. Of course, museums or exhibitions (or towns) are physically organised in some rigid way. When organising the physical layout of the exposition areas a specific criterion for items placement has to be chosen, according to a "default" perspective of information presentation, either chosen by an architect or imposed by geographical constraints. It may happen that the default physical organisation does not meet directly the visitor's expectations, possibly making it difficult to build a personal route. In the case of a museum, paper maps, provided at the entrance, may help visitors navigate through the different rooms to find those items that meet their own interests. Written and tape recorded guides can provide useful description of the objects displayed but they do not guarantee flexibility, either because of technological constraints (i.e. audio tapes force a predefined path) or because the descriptions are not coherently related to each other (e.g., CD-ROMs support free movements but not connected presentations).

On the other hand, virtual museums (implemented for example as hypertextual resources browsable from the World Wide Web) may offer a more flexible (dynamically computed) object display determined by the visitor's individual preferences. The visitors can play around a clickable representation of the museum rooms and objects, getting informative cards on what they are most interested in. However, with purely virtual spaces the visitor may perceive objects differently (e.g., different dimensions and colours) and miss the emotional involvement experienced with real objects (e.g., being in front of Mona Lisa at the Louvre is quite a different experience from looking at its reproduction on the Web).

New hardware technology allows the fruition of virtual repositories of information while enjoying the physical space: for example, kiosks or portable devices may allow the access to a portion of the virtual information space relevant for the object in front of the visitor. More generally, information technology can be exploited to modify the living environment, helping to augment the functionalities of every-day things, therefore enlarging human possibilities in interacting with objects. Figure 1 suggests one possible scenario of this kind, where visitors of a museum or an exhibition may enjoy a personalised tour through the interaction with "augmented objects". Each visitor is equipped with a palmtop computer endowed with headphones, on which an infrared receiver is mounted. Each meaningful physical location has a small (power-autonomous) infrared emitter, sending a code that uniquely identifies it (see fig. 2).

Exploiting the infrared signals, the system is able to identify when the visitor reaches a certain physical location and can activate a relevant portion of the information repository loaded on the palmtop. Meaningful information are selected and organised to be played as audio messages or displayed on the palmtop screen. Adaptive and dynamic hypertext technology is exploited to tailor a presentation according to the visitor interests, the actual context of the visit and so on. In this way, the system guides the visitor through the museum and provides information by means of audio messages: the visitor can get instructions for finding the objects of interest, hear descriptions with references to other interesting objects in the museum, ask for additional information and receive suggestions on next steps. Therefore, he is provided with a personalised guide for exploring the augmented physical space [NPS+97].

We are currently investigating the features of this scenario inside the HyperAudio project, a project IRST is developing in collaboration with the Civic Museum of Rovereto (Italy). The results gained inside HyperAudio will contribute to more advanced research efforts towards a richer interaction scenario, to be explored jointly with other partners inside the HIPS European project [HIP].

In the following sections we describe the critical research issues involved in HyperAudio and the approach that has been adopted for the development of the final prototype. In particular we focus on new opportunities opened by the use of this tool for active visiting.

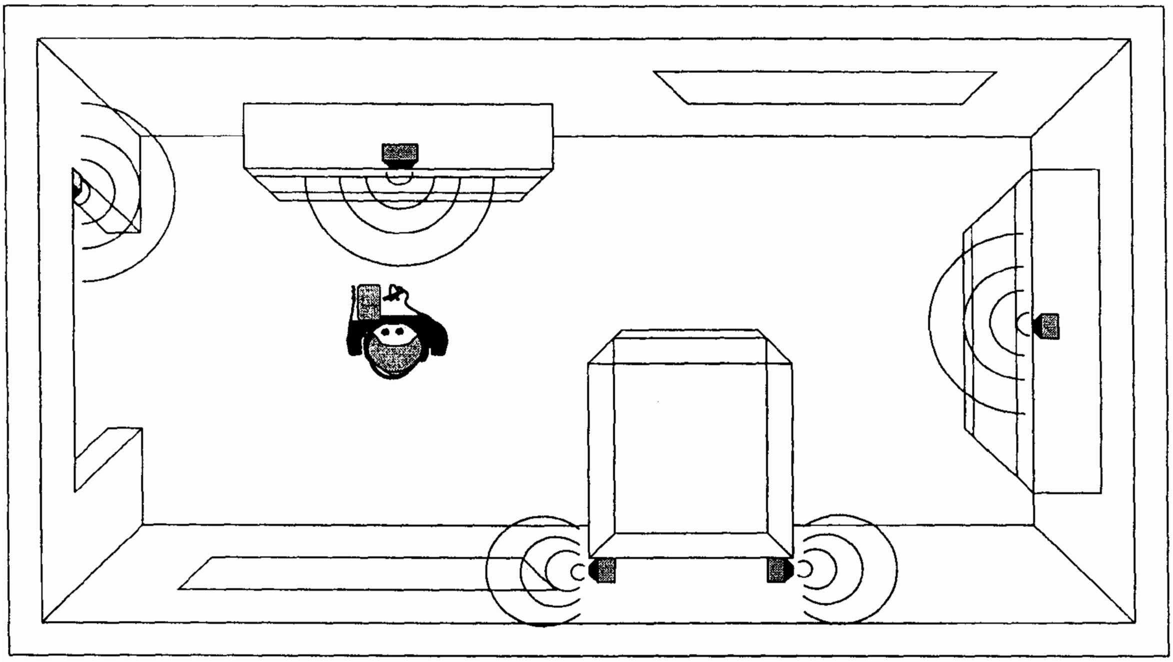

Figure 1: A visitor exploring an augmented room in a museum.

3. Physical and Virtual Museum vs. Augmented Museum

The peculiarity of integrating a virtual information space (e.g. a hypertext) onto a physical space goes beyond the simple juxtaposition of the two spaces; what we get is an *augmented space* with new interaction modalities and new navigation problems:

- Moving around the physical space, approaching meaningful physical locations, the visitor implicitly "clicks" on corresponding points in the information space.
- After getting information about an object (navigation in the information space), the visitor may decide to move in a direction that was explicitly or implicitly suggested by the message as a next interesting step. But the visitor may also decide to ignore the suggested tour thread. Therefore, the navigation in the informative and in the physical space do not necessarily coincide.
- Physical hints may attract the visitor's interest more than the proposed follow-up hypertextual links: he may be distracted by interesting objects close by or he may have personal intuitions about semantic relations between objects that make him stray from the chosen path.

Exploiting the features of the augmented space the system can play a new role in the information-providing task: not only can the system provide the visitor with information tailored to his own interests and interaction history (as it happens with adaptive hypermedia, [cf. Bru96], but it can also support the visitor in his own exploration of the physical space, helping him to find what he is looking for, suggesting new interesting physical locations, trying to avoid his getting lost in the physical space or misunderstanding what he sees.

As an example of how an augmented museum would assist a personalised visit, let us suppose that the visitor is in front of a case containing a stuffed squirrel. The system may provide a description of the animal, suggesting additional explorable information about its predators. If the visitor asks for this type of proposed exploration (by clicking a button on the palmtop screen), the system may provide a description of animals like eagles, foxes,

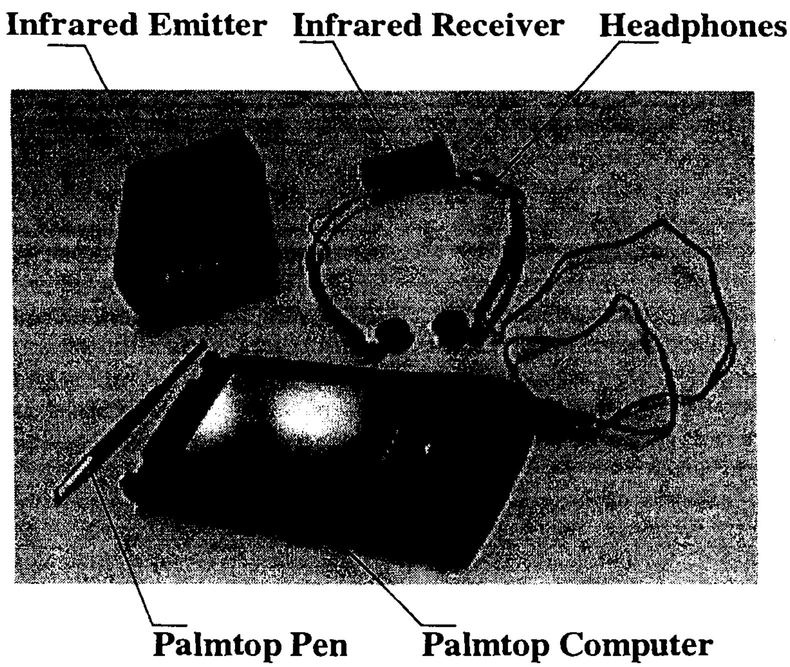

Figure 2: Equipment for a visitor in an augmented space.

etc. If the cases containing the predators are close by, the system could suggest looking at them, explaining how to reach the new location. However, if the shown predators are very far away orthe visitor has displayed interest only in rodents, the system may decide not to propose the site.

Figure 3 illustrates the scenario in which a flexible hyperspace is created taking into consideration both the semantic relations between concepts, the physical constraints as well as the user model and the interaction history.

The nodes in the virtual space are created on the fly in order to adapt to the preferences of the current visitor and to what he has already seen: the message contained in the node created for the current physical location is played as an audio description; additional explorable links to relevant portions of the (dynamic) virtual space are displayed on the palmtop screen.

3.1 The Notion of Path in an Augmented Space

When a physical space is coupled with a virtual space the notion of path is twofold:

- it has a semantic dimension because the purpose of the movement is to get information and not simply to walk around: visiting new points in the space implies that new information is gathered;
- it has the usual spatial dimension in which the visitor moves from one point to another. The spatial dimension itself is twofold:
- the movement can be in the physical space, involving perceptual experiences with real objects and physical strain, or
- the movement can be in the virtual space, i.e. browsing and jumping in the hyperspace.

In an information space there may exist *predefined paths,* i.e. sequences of physical/virtual locations preselected by the designer, that offer particular perspectives on the content. In any case the visitor is not compelled to follow a predefined path, but may take diversions, both in physical and virtual space: we will call the actual sequence of undertaken steps the *actual path*.

Moreover, in an augmented information space, a portable system can suggest at any physical/virtual location next alternative steps, dynamically determined according to the visitor interests, the previous steps and possibly educational strategies. To carry out this active role, the system has also to take into consideration the peculiarities of the two dimensions (semantic and spatial). One key element is the different notion of *distance*: there is a *semantic distance* as a measure of how much two concepts are related to each other and there is a *spatial distance* as a measure of how many points need to be touched before getting to a certain destination. The paths that the system suggests to the visitor should not involve big distances between steps, to limit the mental and physical effort in the navigation and the risk that the visitor gets lost in the physical and information space. To produce coherent information presentations, the system should build object descriptions that involve only "close" semantic concepts. The semantic distance depends on many factors: not only on proximity in an ontological representation of the world but also on contextual saliency [Mil97b]. Therefore the semantic distance of two concepts may vary during the presentation[4]. Distances in the spatial dimension depend instead on the structure of the space. In the physical space the metric used to compute distances should consider the presence of obstacles (walls, stairs, lifts), possibly assigning weights to their negative effect (e.g. stairs involve much more strain that lifts, changing room usually involves less burden than changing floor, and so on). In a virtual space, instead, (e.g. a hypertextual informative space) distances are computed in terms of links you need to traverse. In the optimal situation of a dynamic hypertext in which the relevant hypertextual links available to the user are computed on the fly, distances could be reduced by the system in an automatic way.

4. SYSTEM ARCHITECTURE FOR A PERSON-ORIENTED ELECTRONIC GUIDE

The most suitable hardware for a portable device intended to guide the visitor flexibly throughout an augmented space is a mobile platform of reduced size (e.g. palmtop computer) together with some devices used for position localisation. The more sophisticated the localisation mechanism is (from

Figure 3: How a virtual space is dynamically created in an augmented space

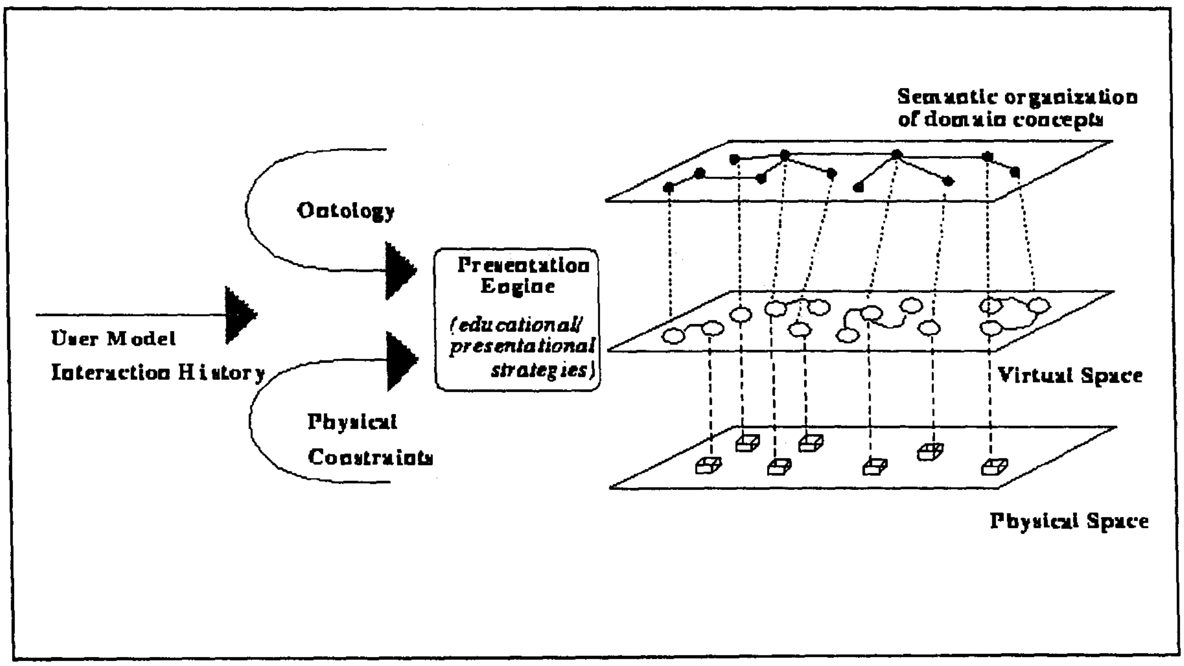

the simple confirmation of the visitor's presence in front of the object, up to the determination of his distance and orientation in the room) the more effective the system guiding can be.

In HyperAudio we are currently using a Newton equipped with an infrared sensor. Small infrared transmitters are placed in meaningful points of the museum rooms, without damage to the building, since no cabling is required (each transmitter has an internal battery for power supply). Each transmitter emits an infrared signal that encodes a number uniquely identifying it. Whenever the palmtop enters in the area of a transmitter, the receiver (located on the headset) catches the identifying code and locates the visitor's position (the system is provided with a map of the transmitters placement). The maximum area affected by a transmitter can reach 10 metres in distance with an angle of 30 degrees. The covered area can be reduced in order to increase the precision of the localisation method when many transmitters may interfere with each other.

4.1 IMPLICIT AND EXPLICIT INPUT

The ability of the system to track the visitor's movements in the physical space allows the system to propose information relevant to the visitor's needs and interests even without an explicit request. Sensors provide an implicit[5] input for system decisions. The visitor is free to look around in the room and focus his attention on the objects displayed: he is not compelled to check on the palmtop screen if and where[6] relevant information can be found. If he stops in front of an interesting object a descriptive audio message is automatically played. It is up to the system to decide if the visitor's movements have to be intended as a relevant input or not. For example, if the system has suggested to go to a given location, sensors activated along the route will not trigger presentations (although the information that the visitor has passed through other locations will be retained for other purposes).

Furthermore, the visitor has also means to explicitly interact with the system, exploiting its suggestions for studying in more depth interesting topics. The palmtop screen displays virtual buttons the person can click on to get additional information that elaborate on the audio message content. For example, buttons can fire comparisons to similar objects (paintings of the same author, paintings of the same period, etc. ...)(cf. [Mil97a] or can activate the presentation of instructions (with maps and verbal directions) on how to reach related interesting objects mentioned in the initial audio message.

4.2 OUTPUT

System output includes either content presentation and next steps suggestion. In our HyperAudio project the content is mainly presented in linguistic form through an audio device. Pre-recorded audio messages, appropriately assembled together, are played either on user request or automatically when the system identifies this is suitable. Museum maps, follow up browsable links and system commands (e.g. to stop the audio play or to show a global map of the museum with the current position of the visitor) are presented as clickable icons or text.

In the following, we will also point out how adaptivity can be enhanced in the output presentation by exploiting the achievements in the field of natural language generation and speech technology.

4.3 MODULES AND KNOWLEDGE SOURCES

The HyperAudio architecture is mainly organised as sketched in figure 3 above. The core of the system is a *presentation composer* that determines the relevant information to be presented: it actually builds the nodes of the virtual space the visitor can (implicitly or explicitly) click onto, giving rise to a dynamic, adaptive virtual space. The nodes are built on the fly as the interaction progresses and contain an audio file to be played, clickable links displayed on the palmtop screen pointing to follow-up content, clickable icons corresponding to interesting following steps in the suggested path. To produce the most effective information presentation, the presentation composer accesses different knowledge sources:

Ontology: An ontology reflecting the semantic organisation of the most relevant domain concepts is maintained to allow the computation of semantic distances and attribute saliency.

Macronodes: Information is organised in a network of macronodes (see fig. 4). Each macronode includes a set of linkable audio files (among which the ones actually played are chosen), pointers to other related macronodes (also expressing the particular rhetorical relation occurring between the macronodes' content[7]), the type of text (e.g. general description, historical information, and so on), a pointer to the associated semantic concept, and possibly a link to the physical locations for which the macronode is relevant.

Museum maps: The system has also a set of maps of the physical space and the capability of limited reasoning about spatial relationships both linguistically and graphically.

User Model: The system maintains a model of the particular user with whom it is currently interacting. Our user model has two kinds of information: what the user has been exposed to or is assumed to know, and what the user seems to be interested in. The user's knowledge model is based on an initialisation (depending on a user profile, e.g. expert or novice) and on a modelization of what the user has become aware of. The user's interest model provides a model of the potential interest of the user and consists on an activation/inhibition weighted network whose nodes are associated with ordered sets of individual concepts.

The user model is exploited during the process of content generation in order to avoid repetitions and false assumptions, and to promote other objects potentially relevant from the user point of view. The updating of the user model depends on both what the user really chose and what he rejected (i.e. stopping an audio play, not clicking a particular choice, unexpectedly changing the physical path, and so on).

Interaction History: The system also maintains a history of the interaction, keeping track of what the visitor has seen, heard and asked for. This information is used to effect further output, for example reminding preceding presentation sequences.

To generate presentation nodes, the composer first chooses the most appropriate subset of audio files from the current macronode and plays them. Secondly, it decides which rhetorical relations to other macronodes are relevant in the present context and displays them as clickable links to related information. Finally, if needed, it integrates the presentation with maps and spatial directions, for example instructions on how to reach a certain location or relevant rooms.

The system adaptivity therefore emerges in different forms, both in the information provided and in the further steps suggested. Adaptivity in the information is not limited to content selection but may also affect the choice of language style: the context of interaction may suggest different linguistic choices, for example the selection of deictic referring expressions (instead of indefinite or definite ones) in case an audio message describing an object is heard when the visitor is in front of the object[8]. Adaptivity in the selection of following steps to suggest consists in the identification of "locally" interesting links, according to user model and interaction:

- new virtual locations to visit are suggested depending on rhetorical structure and text typology;
- new physical locations to visit are suggested depending on ontology relevance and physical constraints.

Adaptivity can be further enhanced when the system pursues more ambitious goals concerning educational strategies and has to negotiate between the visitor's initiative and its own promotion constraints. A simple example of this is that a system attracts the visitor toward a predefined path when this is compatible with the visitor's interests.

5. FURTHER DEVELOPMENTS

In this paper we have presented a system that produces informative guidance through an exposition area by exploiting data coming from the physical world (the movements of the visitor inside the space), the history of the interaction up to the current moment (descriptions already given and requests made) and a repository of information related to the domain. The actual presentation of information is:

- dynamically assembled;
- adaptive, taking into account the user's interests and knowledge;

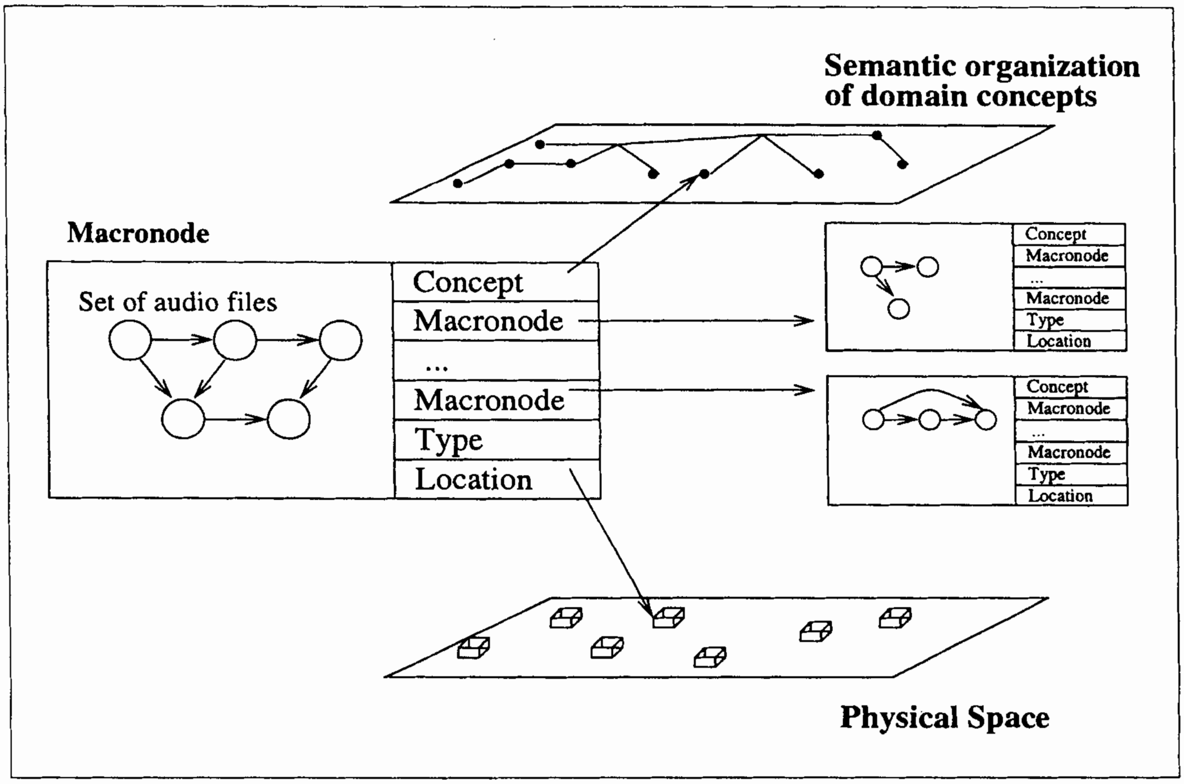

Figure 4: The macronode structure.

- multimodal, integrating audio information with a display on the screen;
- integrated with maps and spatial directions;
- used interactively, with the user allowed to explicitly request for new information.

The solutions adopted in HyperAudio for the accomplishment of these features can be improved in important ways. For example, for the assembling of the audio presentation we adopt the simple solution of selecting and concatenating pre-existing audio files, since our main initial research concern was especially focused on other navigation issues. However, we believe that more sophisticated techniques developed in the field of natural language generation could be profitably exploited in this context. The content of the message, in fact, could be produced with different degrees of complexity, yielding increasingly flexible presentations, for example by:

- modifying existing audio/text files through the introduction of cohesive devices that improve the fluency of the final message (anaphora introduction, introduction of linking sentences that signal comparisons, prosody adjustments, choice of cue-phrases, and so on) [O'D97];
- generating the messages from scratch, with a flexible content selection and organisation, starting from the knowledge represented in the ontology [KMOO96], [MOOK97]. Note that not always this is the most appropriate solution even if you do not consider the technical difficulty: pre-existing texts prepared by critics or domain experts can be good material to exploit.

This and other challenging research issues emerging in the interaction scenario described here are currently being investigated inside HIPS (Hyper-Interaction within the Physical Space), a new large European project under the Esprit programme Intelligent Information Interfaces (I^3) [HIP][9].

NOTES

1. ILEX (Intelligent Labelling EXplorer) generates personalised descriptions of artifacts in a (virtual) gallery adapting the information conveyed according to the features of the particular visitor and to the interaction history.
2. PEBA-II takes an underlying knowledge base containing information on the animal kingdom and produces from it adaptive textual descriptions for the animals similar in content to encyclopedia entries.
3. By 'static' here we mean that the content of the catalogue is a-priori fixed, even though the user is free to browse around it.
4. For example, when describing the lion -a carnivorous animal- we should avoid describing the flora of its living environment, unless this is motivated by the particular user interests.
5. Here *implicit* means that there is no intentionality on the part of the visitor to communicate his position to the system.
6. If the system automatically computes the relevant portion of the information repository to activate according to the visitor's position, the visitor is not compelled to browse through the information in search for what he is looking at.
7. For the modeling of the rhetorical relations expressing how text spans coherently relate to each other we have adopted the well known Rhetorical Structure Theory (RST), widely used in many generation systems, cf. [MT87].
8. In the simplest case, alternative audio messages may be available presenting the same information with different language style.
9. The HIPS consortium includes: University of Siena (coordinating partner), CB&J (France), GMD (Germany), IRST (Italy), SIETTE-Alcatel (Italy), SINTEF (Norway), University of Dublin and University of Edinburgh.

REFERENCES

[Bru96] Peter Brusilovsky. Methods and techniques of adaptive hypermedia. *User Modeling and User Adapted Interaction*, 6(2-3), 1996.

[DM96] Robert Dale and Maria Milosavljevic. Authoring on demand: Natural language generation in hypertext documents. In *Proceedings of the First Australian Document Computing Conference*, Melbourne, Australia, March 1996.

[Fle97] *Proceedings of the Flexible Hypertext Workshop*, Southampton, UK, April 1997. A Workshop held in conjunction with the Eight ACM International Hypertext Conference (Hypertext97).

[HA] WWW home page for the HyperAudio project:

http://ecate.itc.it:1024/projects/hyperaudio.html

[HIP] WWW home page for the HIPS project: http://www.ing.unisi.it/lab_tel/hips/hips.html.

[KMOO96] Alistair Knott, Chris Mellish, Jon Oberlander, and Mick O'Donnell. Sources of flexibility in dynamic hypertext generation. In *Proceedings of the Eight International Workshop of Natural Language Generation*, Herstmonceaux Castle, Sussex, UK, June 1996.

[Mil97a] Maria Milosavljevic. Augmenting the user's knowledge via comparison. In *Proceedings of the 6th International Conference on User Modelling*, Sardinia, Italy, 1997.

[Mil97b] Maria Milosavljevic. Content selection in comparison generation. In *Proceedings of the Sixth European Workshop on Natural Language Generation*, Duisburg, Germany, 1997.

[MOOK97] C. Mellish, J. Oberlander, M. O'Donnell, and A. Knott. Exploring a gallery with intelligent labels. In *Proceedings of the Fourth International Conference on Hypermedia and Interactivity in Museums* (ICHIM97), Paris, September 1997.

[MT87] William C. Mann and Sandra A. Thompson. Rhetorical Structure Theory: A Theory of Text

Organization. In Livia Polanyi, editor, *The Structure of Discourse*. Ablex Publishing Corporation, 1987.

[NPS+97] Elena Not, Daniela Petrelli, Oliviero Stock, Carlo Strapparava, and Massimo Zancanaro. Augmented space: Bringing the physical dimension into play. In *Proceedings of the Flexible Hypertext Workshop* [Fle97]. Also available as IRST Technical Report #9702-08.

[O'D97] Mick O'Donnell. Variable-length on-line document generation. In *Proceedings of the Flexible Hypertext Workshop* [Fle97].

[ST93] O. Stock and The ALFRESCO Project Team. ALFRESCO: Enjoying the combination of NLP and hypermedia for information exploration. In M.T. Maybury, editor, Intelligent Multimedia Interfaces. AAAIPress/MIT Press, Menlo Park CA/Cambridge MA, 1993.

[Sto94] Oliviero Stock. Natural language in multimodal human-computer interfaces. IEEE Expert. Intelligent Systems and Their Applications, 9(2), April 1994.