DOCUMENT RESUME

ED 482 094   IR 058 782

AUTHOR   Sparacino, Flavia
TITLE   The Museum Wearable: Real-Time Sensor-Driven Understanding of
              Visitors' Interests for Personalized Visually-Augmented
              Museum Experiences.
PUB DATE   2002-04-00
NOTE   33p.; In: Museums and the Web 2002: Selected Papers from an
              International Conference (6th, Boston, MA, April 17-20,
              2002); see IR 058 778.
AVAILABLE FROM For full text: http://www.archimuse.com/mw2002/
              papers/sparacino/sparacino.html/.
PUB TYPE   Reports - Descriptive (141) -- Speeches/Meeting Papers (150)
EDRS PRICE   EDRS Price MF01/PC02 Plus Postage.
DESCRIPTORS   *Audiovisual Aids; *Guides; Information Technology; *Museums
IDENTIFIERS   Handheld Computers; Mobile Equipment

ABSTRACT
              This paper describes the museum wearable: a wearable computer
that orchestrates an audiovisual narration as a function of the visitors'
interests gathered from their physical path in the museum and length of
stops. The wearable consists of a lightweight and small computer that people
carry inside a shoulder pack. It offers an audiovisual augmentation of the
surrounding environment using a small, lightweight eyepiece display (often
called private-eye) attached to conventional headphones. Using custom-built
infrared location sensors distributed in the museum space, and statistical
mathematical modeling, the museum wearable builds a progressively refined
user model and uses it to deliver a personalized audiovisual narration to the
visitor. This device enriches and personalizes the museum visit as a visual
and auditory storyteller that is able to adapt its story to the audience's
interests and guide the public through the path of the exhibit. (Contains 19
references.) (Author/MES)

PAPERS
Museums and the Web 2002 [1]

The Museum Wearable:

Real-Time Sensor-Driven Understanding Of Visitors' Interests For Personalized Visually-Augmented Museum Experiences

Flavia Sparacino, MIT Media Lab, USA

http://www.media.mit.edu/~flavia/projects.html

Register
Workshops
Sessions
Speakers
Interactions
Demonstrations
Exhibits
Events
Best of the Web
Key Dates
Boston
Sponsors

A&MI

Search
A&MI

Join our Mailing List.
Privacy.

Abstract

This paper describes the museum wearable: a wearable computer
which orchestrates an audiovisual narration as a function of the
visitors' interests gathered from their physical path in the museum and
length of stops. The wearable consists of a lightweight and small
computer that people carry inside a shoulder pack. It offers an
audiovisual augmentation of the surrounding environment using a
small, lightweight eye-piece display (often called private-eye) attached
to conventional headphones. Using custom built infrared location
sensors distributed in the museum space, and statistical mathematical
modeling, the museum wearable builds a progressively refined user
model and uses it to deliver a personalized audiovisual narration to the
visitor. This device will enrich and personalize the museum visit as a
visual and auditory storyteller that is able to adapt its story to the
audience's interests and guide the public through the path of the
exhibit.

Keywords: wearable computer, handheld, personalized experience,
Bayesian networks, location aware mobile computing

1. Introduction

In the last decade museums have been drawn into the orbit of the leisure
industry and compete with other popular entertainment venues, such as
cinemas or the theater, to attract families, tourists, children, students,
specialists, or passersby in search of alternative and instructive entertaining
experiences. Some people may go to the museum for mere curiosity,
whereas others may be driven by the desire for a cultural experience. The
museum visit can be an occasion for a social outing, or become an
opportunity to meet new friends. While it is not possible to design an exhibit
for all these categories of visitors, it is desirable to offer exhibit designers
ways to attract as many as possible amongst the variety of individual visitors
or visitor categories. Technology can help offer the museum public
opportunities to personalize their visits according to their desires and
expectations.

Traditional storytelling aids for museums have been signs and text labels,
spread across the exhibit space; exhibit catalogues, typically sold at the
museum store; guided tours, offered to groups or individuals; audio tours;
and more recently video or multimedia kiosks with background information
on the displayed objects. Each of these storytelling aids has advantages and

disadvantages. Catalogues are usually attractive and well done, yet they are often too cumbersome to carry around during the visit as a means to offer guidance and explanations. Guided tours take away from visitors the choice of what they wish to see and for how long. They can be highly disruptive for the surrounding visitors, and their effectiveness strictly depends on the knowledge, competence, and communicative skills of the guide. Audio tours are a first step to help augment the visitor's knowledge. Yet when they are button activated, as opposed to having a location identification system, they can be distracting for the visitor. The information conveyed is also limited by the only-audio medium: it is not possible to compare the artwork described with previous relevant production of the author, nor to show other relevant images. Interactive kiosks are more frequently found today in museum galleries. Yet they are physically distant from the work they describe, thus do not support the opportunity for the visitors to see, compare, and verify the information received against the actual object. The author's experience suggests that when extensive Web sites are made available through interactive kiosks placed along the museum galleries, these may absorb lengthy amounts of the visitors' museum time, thereby detracting from, rather than enriching, the objects on display. Finally, panels and labels with text placed along the visitors' path can interrupt the pace of the experience as they require a shift of attention from observing and contemplating to reading and understanding (Klein, 1986).

Indeed, when we walk through a museum there are so many different stories we could be told. Some of these are biographical about the author of an artwork; some are historical and allow us to comprehend the style or origin of the work; and some are specific about the artwork itself, in relation to other artistic movements. Museums usually have large Web sites with multiple links to text, photographs, and movie clips to describe their exhibits. Yet it would take hours for a visitor to explore all the information in a kiosk, to view the VHS cassette tape associated with the exhibit, and read the accompanying catalogue. Most people do not have the time to devote nor motivation to assimilate this type of information, and therefore the visit to a museum is often remembered as a collage of first impressions produced by the prominent feature of the exhibits, and the learning opportunity is missed. How can we tailor content to the visitors in a museum, during their visits, to enrich both the learning and entertainment experience ? We want a system which can be personalized to be able to dynamically create and update paths through a large database of content and deliver to the users in real time during the visit all the information they desire. If the visitors spend a lot of time looking at a Monet, the system needs to infer that the users likes Monet, and update the paths through the content to take that into account. This research proposes the museum wearable as an effective way to turn this scenario into reality.

This document illustrates the hardware, authoring techniques, and software created for the construction of the museum wearable. It speculates on its ability to facilitate a new style of exhibit design, and postponesits, as future work, the assessment at the exhibit's site, of the museum wearable's contribution to the public's experience.

2. The Museum Wearable

Wearable computers have risen to the attention of technological and scientific investigation (Starner, 1997) and offer an opportunity to "augment" for visitors their perception/memory/experience of the exhibit in a personalized way. The museum wearable is a wearable computer which orchestrates an audiovisual narration as a function of the visitors' interests

3

gathered from their physical path in the museum and length of stops. It offers a new type of entertaining and informative museum experience, more similar to mobile immersive cinema than to the traditional museum experience [figure 1].

The museum wearable is made by a lightweight CPU hosted inside a small shoulder pack and a small, lightweight private-eye display. The display is a commercial monocular, VGA-resolution, color, clip-on screen attached to a pair of sturdy headphones. When wearing the display, after a few seconds of adaptation, the user's brain assembles the real world's image, seen by the unencumbered eye, with the display's image seen by the other eye, into a fused augmented reality image [figures 2, 3].

The wearable relies on a custom-designed long-range infrared location-identification sensor to gather information on where and how long the visitor stops in the museum galleries. A custom system had to be built for this project to overcome limitations of commercially available infrared location identification systems such as short range and narrow cone emission. The location system is made by a network of small infrared devices, which transmit a location identification code to the receiver worn by the user and attached to the display glasses. The transmitters have the size of a 9V battery, and are placed inside the museum, next to the regular museum lights. They are built around a microcontroller, and their signal can be reliably detected at least as far as 30 feet away within a large cone range of approximately ten to thirty degrees, according to the area that needs to be covered, and up to 100 feet along a straight line. The emitter location identification tags have been embedded inside standard light fixtures to allow the exhibit designer to easily place them in the museum, next to the regular museum lights, and using the same power rack as the regular museum spotlights [figures 4, 5, 6, 7].
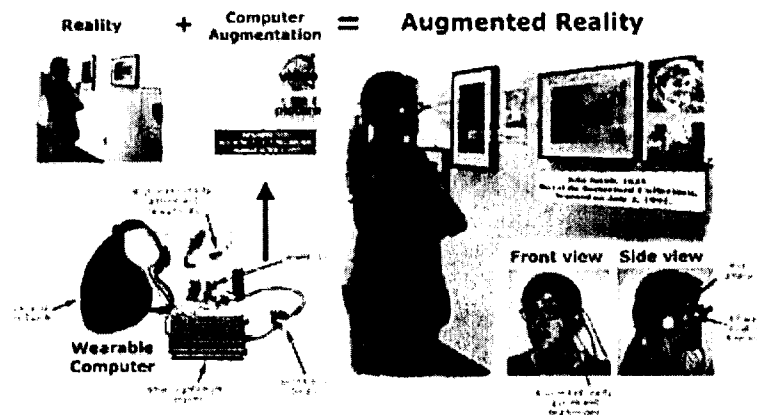

Figure 1. The Museum Wearable: explanation of concept and application
(click for a detailed image)

4

Figures 2, 3. Camera "wearing" the head mounted display: shows how the user's brain assembles the real world's image seen by the unencumbered eye with the display's image seen by the other eye, into a fused augmented reality image.

Figures 4, 5, 6, 7. Location sensor: emitter tags embedded inside light fixtures

In view of having a museum wearable which can later be expanded to include other sensors, and process information not just from the infrared location sensor, but for example also from a small camera processing images in real time, a commercially available processing unit has been selected for this project. The processor of choice is a small-sized laptop computer: the SONY picturebook PCG-C1VPK, selected for its combined size, weight, computing power, multimedia capabilities, and long-lasting batteries. Given that the images generated by the laptop are viewed uniquely through the head mounted display, the LCD screen has been removed from the picturebook, to reduce weight and size. The picturebook features a Crusoe™ processor TM5600 clocked at 667 MHz, and without the LCD weighs only approximately one lb, and has a size of 0.5" X 9.8" X 6.0" (H x W x D). The picturebook has a 15GB capacity hard drive, which allows the programmer to store on the local hard drive several hours (8-10) of MPEG-compressed VGA resolution video (640x480) (approximately one hour of MPEG-compressed 640x480 video per one GB of available space on the internal hard drive). It also has 128 MB SDRAM, which allows the computer to play smoothly audio and video clips, as well as process images in real time when the computer is connected to a camera. The external ports include one USB port which is connected to the infrared receiver with a USB to serial converter, and a VGA and headphone output, which are connected to the video/audio inputs of the head-mounted display of the museum wearable. It also supports one type II card, which can be used to host a PCMCIA card for wireless communication over the Internet or a PCMCIA video acquisition card. All these features, combined with a battery life of 2.5-5.5 hours with the standard lightweight battery, large enough for a single museum visit, make the picturebook an ideal choice for the selected application [diagram 1].

An alternative to the picturebook is the smaller handheld IPAQ pocket PC 3670 [figures 8, 9]. The iPAQ 3670 features 64 MB of SDRAM and a 206-MHz Intel StrongARM SA-1110 32-bit RISC Processor. It has USB or serial connectivity that would interface with the infrared receiver of the museum wearable and it is only 5.11" x 3.28" x 0.62" (HxWxD) in size. The iPAQ 3670 is a desirable solution for the deployment of several museum wearables – which need to work only with the location sensor – in a museum. The iPAQ solution is cost effective because its price, with the necessary accessories, is about half the price of the SONY picturebook, and its size is also smaller.

The size and weight of both the wearable's processor and augmented reality display are critical for a museum application. The display cannot have a heavy and power hungry powering unit which requires frequent battery changes. Glasses also need to easily fit various people's head sizes, with annexed hair style, which is not an easy task. The wearable would be handed out to between ten and one hundred people a day, and therefore needs to be of robust assembly and easy to wear. Two different design solutions were implemented after a thorough iterative design process which included considerations about video resolution, power consumption, purchase availability and foremost, ergonomics, with the inclusion of

adaptability with sensors, weight, and size. The common fashion display features an augmented reality display, which joins together a lightweight VGA resolution color display from the MicroOptical corporation, and a commercial high quality sturdy set of headphones [figure 10]. The high fashion display is a provocative stylish mount, mainly intended for visitors with a strong sense of aesthetics, and suitable for use in wearable fashion shows, to promote a non-nerdy and highly fashionable wear of augmented reality displays. With this design the MicroOptical augmented reality display is rigidly mounted to a pair of Oakley "over the top" glasses, as illustrated in figure 11.


Figure 8. PAQ pocket PC 3670

Figure 9. Sony picturebook with removed LCD

file://E:\MW2002\papers\sparacino\sparacino.html   5/22/2003

Figures 10,11. The common fashion and high fashion private eye display

Summary of elements of the museum wearable hardware:

- containing shoulder pack
- computer (CPU): SONY picturebook from which the display has been removed to reduce weight
- Head Mounted Display (HMD): VGA resolution MicroOptical clip-on mounted on sturdy headphones with a custom mount
- HMD's powering unit: hosted inside the containing shoulder pack
- Infrared receiver: the sensor is located on top of the headphones and the receiver circuit is located inside the containing shoulder pack.

9

HARDWARE

```
┌─────────┐        ┌──────────┐  SERIAL ┌──────────┐ VGA ┌────────────────────┐
│IR EMITTER│  ≈≈    │IR        ├─────────┤PENTIUM II ├─────┤ HEAD MOUNT DISPLAY │
│TAG      │        │RECEIVER  │         │400 MHz   │     └────────────────────┘
└─────────┘        └──────────┘         └──────────┘
```

Diagram 1: The museum wearable hardware architecture

The museum wearable plays an interactive audiovisual documentary about the displayed artwork on the private-eye display. Each mini-documentary is made of small segments which vary in size from twenty seconds to one and half minutes. A video server, written in C++ and DirectX-8, plays these assembled clips and receives TCP/IP messages from another program containing the information measured by the location ID sensors. This server-client architecture allows the programmer to easily add other client programs to the application, to communicate to the server information from other possible sources such as sensors or cameras placed along the museum aisles to measure how crowded the galleries are or how often a certain object has been visited. The client program reads IR data from the serial port, and the server program does inference, content selection, and content display using DirectX for full screen play back of the MPEG compressed clips [diagram 2].

3. Visitors' interests and visitor types

To identify visitor preferences, museums seek to identify target groups: individuals who share common traits such as culture, ethnic or social affiliation, educational level, and leisure preferences. Curators and designers need to assess the basic knowledge and expectations of these groups to be able to reach, communicate with, and stimulate curiosity in all of their visitors. Eleanor Hooper-Greenhill identifies target groups which include families, school parties, other organized educational groups, leisure learners, tourists, the elderly, and people with visual, auditory, mobility or learning disabilities (Hooper-Greenhill, 1999, p. 86). She then suggests a partition of museum resources, to target, attract and entertain these different groups. During a personal interview, Beryl Rosenthal, director of exhibitions at the MIT Museum, described a more sophisticated visitor type classification. She identified stroller moms, accompanied by children three years old or younger; window shoppers: families who cruise through the museum in search of an alternative leisure experience; button pushers, typically adolescents; school groups; the date crowd; and the PhDs, who want to know (and criticize) everything in the museum. Young visitors, children 5-14 also represent a separate group of visitors with different learning needs and curiosities than the other groups. While this colorful classification well depicts the variety of public that museums need to attract, entice and educate, it is too sophisticated to model mathematically, at least initially.
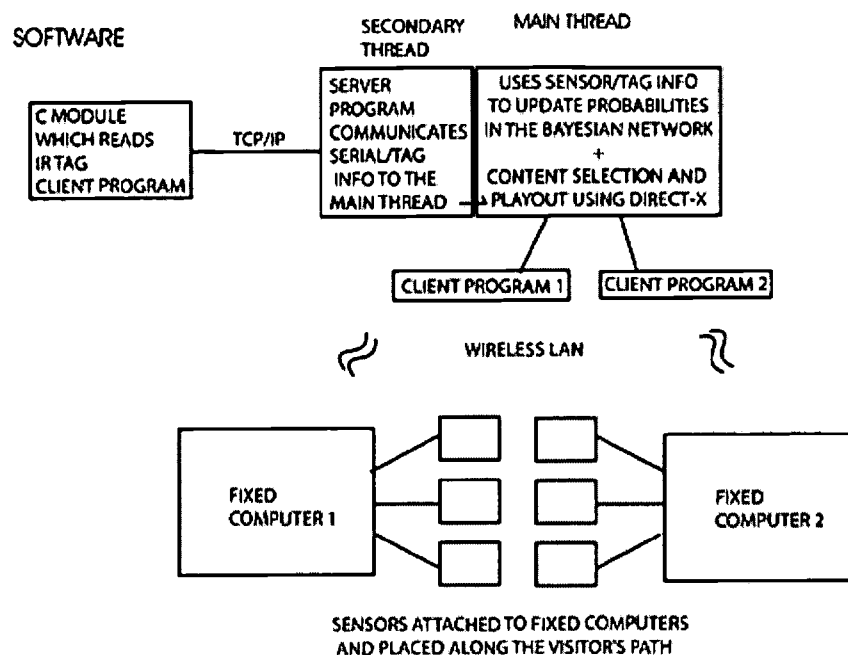
Diagram 2: The museum wearable software architecture

More usefully for this research, Dean generalizes museum visitors in three broad and much simpler categories (Dean, 1994, pp. 25-26). The first category includes what he calls the "casual visitors": people who move through a gallery quickly and who do not become heavily involved in what they see. Casual visitors use some of their leisure time in museums but do not have a strong stimulus or motivation to deepen their knowledge about the objects on display. The second group, the "cursory visitors," show instead a more genuine interest in the museum experience and their collections. According to Dean, these visitors respond strongly to specific objects that stimulate their curiosity and wander through the gallery in search of further such stimulus for a closer exploration of the targeted objects. They do not read every label nor absorb all available information, but will occasionally read and spend time in selected areas or with selected objects of interest they encounter in the galleries. The third group is a minority of visitors who thoroughly examine exhibitions with much more detail and attention. They are learners who will spend an abundance of time in galleries, read the text and labels, and closely examine the objects. Dean attributes differences between "people who rush", "people who stroll", and "people who study" to different prior experiences and educational level. Yet he states that it is important for museums to be equipped to communicate with and interest all visitors by scaling and designing an exhibit so that it offers entertainment to the "stroller" as well as an opportunity to deepen knowledge for the "learners".

Serrell (1996) also divides visitors into three types: the transient, the sampler and the methodical viewers. She notes that currently museum evaluators are using terms like "streakers, studiers, browsers, grazers and discoverers" to characterize museum visitors' styles of looking at exhibits. But she concludes that this type of categorization is not useful for summative evaluation, suggesting that it is a subjective method of classification, and that it is not fruitful to try to create exhibitions that serve these different styles of visiting. Instead she suggests that a more objective means of classification be found, such as average time spent in the exhibition space.

In accordance with the simplified museum visitor typology suggested by Dean and Serrell, the museum wearable identifies three main visitor types. To offer a more intuitive understanding of these types they have been renamed: the busy, selective, and greedy visitor type. The greedy type wants to know and see as much as possible, and does not have a time constraint; the busy type just wants to get an overview of the principal items in the exhibit, and see little of everything; and the selective type wants to see and know in depth only about a few preferred items. The identification of other visitor types or subtypes has been postponed to future improvements and developments of this research.

The visitor type estimation is obtained probabilistically with a Bayesian network using as input the information provided by the location identification sensors on where and how long the visitor stops, as if the system were an invisible storyteller following the visitor in the galleries and trying to guess preferences based on observation of external behavior.

4. Sto(ry)chastics: Sensor-driven Understanding of Visitors' Interests with Bayesian Networks

In order to deliver a dynamically changing and personalized content presentation with the museum wearable, a new content authoring technique had to be designed and implemented. This called for an alternative method than the traditional complex centralized interactive entertainment systems which simply read sensor inputs and map them to actions on the screen. Interactive storytelling with such one-to-one mappings leads to complicated control programs which have to do an accounting of all the available content, where it is located on the display, and what needs to happen when/if/unless. These systems rigidly define the interaction modality with the public, as a consequence of their internal architecture, and lead to presentations which have little depth of content, are hard to modify, ad hoc, and prone to error. The main problem with such content authoring approaches is that they acquire high complexity when drawing content from a large database, and that once built, they are hard to modify or to expand upon. In addition, when they are sensor-driven they become depended on the noisy sensor measurements, which can lead to errors and misinterpretation of the user input (Sparacino, 1999). Rather than directly mapping inputs to outputs, the system should be able to "understand the user" and to produce an output based on the interpretation of the user's intention in context.

To overcome the limitations of one-to-one mapping systems, the mathematical modeling approach developed to author content for the museum wearable is a real-time sensor-driven stochastic modeling of story and user-story interaction. It has therefore been called sto(ry)chastics. Sto(ry)chastics uses graphical probabilistic modeling of story fragments and participant input, gathered from sensors, to tell a story to the user, as a function of estimated intentions and desires during interaction. With this approach, the coarse and noisy sensor inputs – path and length of stops given by the location sensors – are coupled to digital media outputs via a user model, estimated probabilistically by a Bayesian network.

A Bayesian network is a graphical model which encodes probabilistic relationships amongst variables of interest (Pearl, 1988; Jordan, 1999). Such graphs not only provide an attractive means for modeling and communicating complex structures, but also form the basis for efficient algorithms, both for propagating evidence and for learning about parameters. Bayesian networks encode qualitative influences between variables in addition to the numerical parameters of the probability distribution. As such

they provide an ideal form for combining prior knowledge and data. By using graphs, not only does it become easy to encode the probability independence relations amongst variables of the network, but it is also easy to communicate and explain what the network attempts to model (Smyth, 1997). Graphs are easy for humans to read, and they help focus attention, for example in a group working together to build a model. This allows the digital architect, or the engineer, to communicate on the same ground (the graph of the model) as the curator and therefore to be able to encapsulate the curator's domain knowledge in the network, together with the sensor data.

Sto(ry)chastics uses a Bayesian network to estimate the user's preferences taking the location identification sensor data as the input or observations of the network. The user model is progressively refined as the visitor progresses along the museum galleries: the model is more accurate as it gathers more observations about the user. Figure 12 shows the Bayesian network for visitor estimation, limited to three museum objects (so that the figure can fit in the document), selected from a variety of possible models designed and evaluated for this research. Figures 13 and 14 show state values for the network after the visitor has made a long stop at the first object, followed by another long stop at the second object.
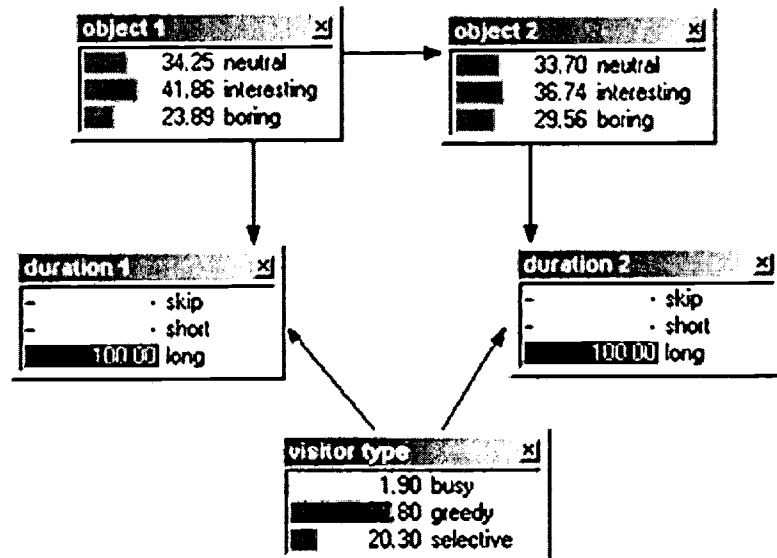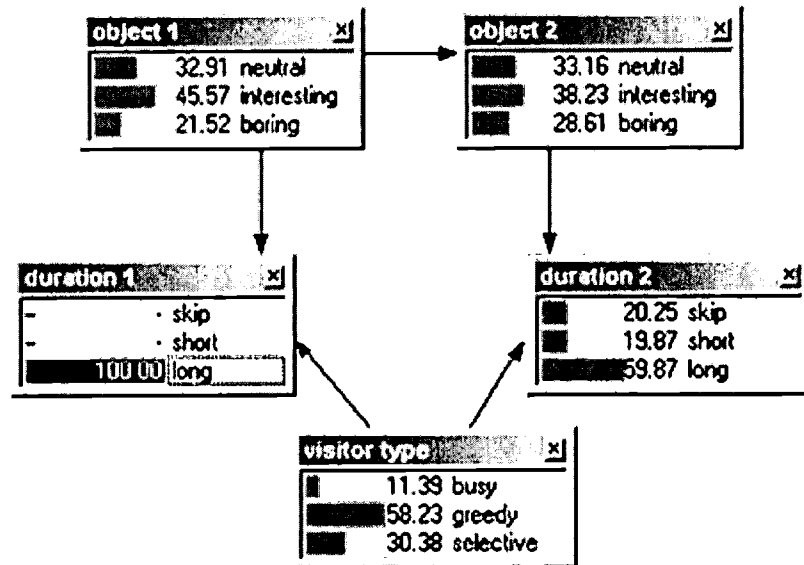

Figure 12. Topology of the Bayesian network for visitor time estimation, limited to three objects/time slices.

The sto(ry)chastics approach has several advantages. It is:

1. Flexible: it is possible to easily test many different scenarios by simply changing the parameters of the system.

2. Reconfigurable: it is also quite easy to add or remove nodes and/or edges from the network without having to "start all over again" and specify again all the parameters of the network from scratch. This is a considerable and important advantage with respect to hard coded or heuristic approaches to user modeling and content selection. Only the parameters of the new nodes and the nodes corresponding to the new

links need to be given. The system is extensible story-wise and sensor-wise. These two properties: flexibility and ease of model reconfiguration, allow the system engineer, the content designer, and the exhibit curator to work together and easily and cheaply try out various solutions and possibilities until they converge on a model which satisfies all the requirements and constraints for their project. A network can also rapidly be reconfigured for another exhibit.

3. Robust: Probabilistic modeling allows the system to achieve robustness with respect to the coarse and noisy sensor data.

object 1
	32.91 neutral
	45.57 interesting
	21.52 boring

object 2
	33.16 neutral
	38.23 interesting
	28.61 boring

duration 1
	· skip
	· short
100.00	long

duration 2
	20.25 skip
	19.87 short
	59.87 long

visitor type
	11.39 busy
	58.23 greedy
	30.38 selective

object 1
	34.25 neutral
	41.86 interesting
	23.89 boring

object 2
	33.70 neutral
	36.74 interesting
	29.56 boring

duration 1
	· skip
	· short
100.00	long

duration 2
	· skip
	· short
100.00	long

visitor type
	1.90 busy
	.80 greedy
	20.30 selective

Figures 13, 14. Values of the network nodes after the visitor has made a long stop at the first object followed by another long stop at the second object.

14

4. Adaptive: sto(ry)chastics is adaptive in two ways: it adapts both to individual users and to the ensemble of visitors of a particular exhibit. For individuals, even if the visitor exhibits an initial "greedy" behavior, it can later adapt to the visitor's change of behavior. It is important to notice that, reasonably and appropriately, the system "changes its mind" about the user type with some inertia: i.e. it will initially lower the probability for a greedy type until other types gain probability. Sto(ry) chastics can also adapt to the collective body of its users. If a count of busy/greedy/selective visitors is kept for the exhibit, these numbers can later become priors of the corresponding nodes of the network, thereby causing the entire exhibit to adapt to the collective body of its users through time. This feature can be seen as "collective intelligence" for an exhibit which can adapt not just to the individual visitors but also to the set of its visitors.

5. Context-sensitive: for any system to be robust and to provide relevant information to its user, it is important to model the context of interaction together with the other system parameters. For example sto(ry)chastics could provide an explanation of a visitor's change of behavior at the museum. If suddenly a greedy type starts making short stops, the system, before concluding that that visitor is actually a selective or busy type, could test if the current time is near closing time for the museum galleries, or if, by use of other room sensors, there is a great crowd in the galleries where the visitor is making short stops. Coming up with the right conclusions, given this type of external information, means that the system is context-sensitive.

Sto(ry)chastics has therefore implications both for the human author (designer/curator) who is given a flexible modeling tool to organize, select, and deliver the story material, as well as the audience, who receives personalized content only when and where it is appropriate.

5. Experimentation: the Museum Wearable at the MIT Museum's Robots and Beyond Exhibit

The ongoing robotics exhibit at the MIT Museum provided an excellent platform for experimentation and testing with the museum wearable [figures 15, 16]. This exhibit, called Robots and Beyond, and curated by Janis Sacco and Beryl Rosenthal, features landmarks of MIT's contribution to the field of robotics and Artificial Intelligence. The exhibit is organized in five sections: Introduction, Sensing, Moving, Socializing, and Reasoning and Learning, each including robots, a video station, and posters with text and photographs which narrate the history of robotics at MIT. There is also a large general purpose video station with large benches for people to have a seated stop and watch a PBS documentary featuring robotics research from various academic institutions in the country.

15

Figures 15 and 16. Visitor with the Museum Wearable at MIT Museum's Robots and Beyond Exhibit

In order to set the initial values of the parameters of the Bayesian network, experimental data was gathered on the visitors' behavior at the Robots and Beyond exhibit. According to the VSA (Visitor Studies Association, http://museum.cl.msu.edu/vsa), timing and tracking observations of visitors are often used to provide an objective and quantitative account of how visitors behave and react to exhibition components. This type of observational data suggests the range of visitor behaviors occurring in an exhibition, and indicates which components attract, as well as hold, visitors' attention (in the case of a complete exhibit evaluation this data is usually accompanied by interviews with visitors, before and after the visit).

During the course of several days a team of collaborators tracked and make annotations about the visitors at the MIT Museum. Each member of the tracking team had a map and a stop watch. Their task was to draw on the map the path of individual visitors, and annotate the locations at which

16

visitors stopped, the object they were observing, and for how long they stopped. In addition to the tracking information, the team of evaluators was asked to assign a label to the overall behavior of the visitor, according to the three visitor categories earlier described: "busy", "greedy", or "selective" [figure 17].
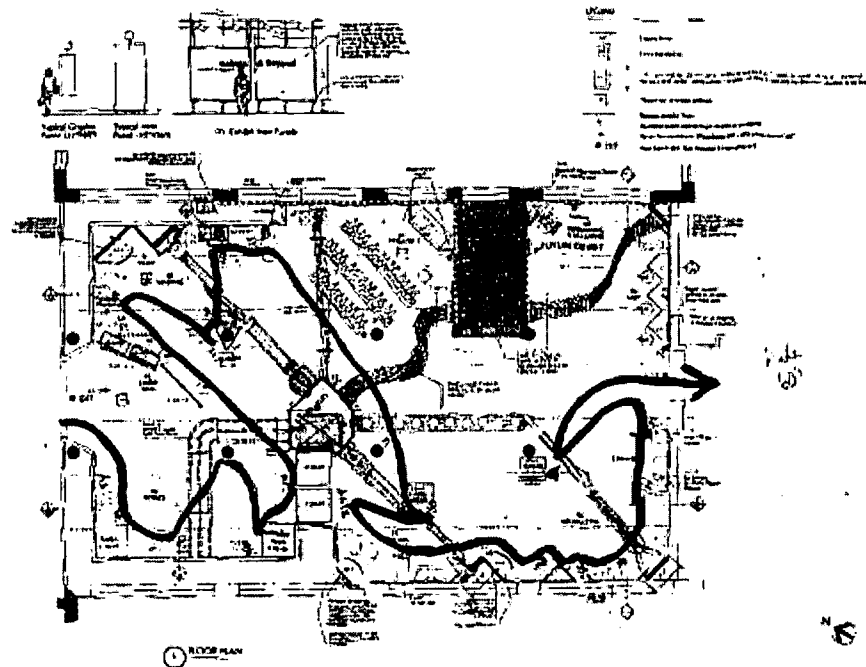

Figure 17. Example of annotations of visitor's path and stop duration at MIT Museum's Robots and Beyond exhibit

A subset of twelve representative objects of the Robots and Beyond exhibit were selected to evaluate this research, to shorten editing time. The geography of the exhibit needed to be reflected into the topology of the network, as shown in figure 18. Additional objects/nodes of the modeling network can be added later for an actual large scale installation and further revisions of this research.

The visitor tracking data is used to learn the parameters of the Bayesian network. The model can later be refined; that is, the parameters can be fine-tuned, as more visitors experience the exhibit with the museum wearable. The network has been tested and validated on this observed visitor tracking data by parameter learning using the Expectation Maximization (EM) algorithm, and by performance analysis of the model with the learned parameters, with a recognition rate of 0.987 (Sparacino, 2001).

The system works in two steps. The first is user type estimation as described above. The next step is to assemble a mini-story for the visitor, relative to the nearest object [figures 19, 20, 21]. Most of the audio-visual material available for use by the museum wearable tends to fall under a set of characterizing topics, which typically define art and science documentaries. This same approach to documentary as a composition of segments belonging to different themes has been developed by Houbart in her work which edits a documentary based on the viewer's theme preferences, as an offline process (Houbart, 1994). The difference between Houbard's work and what the museum wearable does is that the museum wearable performs editing in

17

real time, using sensor input and Bayesian network modeling to figure out the user's preferences (type). After an overview of the audio-visual material available at MIT's Robots and Beyond exhibit, the following content labels, or bins, were identified to classify the component video clips:

- Description of the artwork: what it is, when it was created (answers: when, where, what)
- Biography of author: anecdotes, important people in artist's life (answers: who)
- History of the artwork: previous relevant work of the artist
- Context: historical, what is happening in the world at the time of creation
- Process: particular techniques used or invented to create the artwork (answers: how)
- Principle: philosophy or school of thought the author believes in when creating the artwork (answers: why)
- Form and Function: relevant style, form and function which contribute to explain the artwork.
- Relationships: how is the artwork related to other artwork on display
- Impact: the critics' and the public's reaction to the artwork

This project required a great amount of editing to be done by hand (non automatically) in order to segment the two hours of video material available for the Robots and Beyond Exhibit at the MIT museum in the smallest possible complete segments. After this phase, all the component video clips were given a name, their length in seconds was recorded into the system, and they were also classified according to the list of bins described above. The classification was done probabilistically; that is, each clip has been assigned a probability (a value between zero and one) of belonging to a story category. The sum of such probabilities for each clip needs to be one. The result of the clip classification procedure, for a subset of available clips, is shown in table 1.

To perform content selection, "conditioned" on the knowledge of the visitor type, the system needs to be given a list of available clips and the criteria for selection. There are two competing criteria: one is given by the total length of the edited story for each object, and the other is given by the ordering of the selected clips. The order of story segments guarantees that the curator's message is correctly passed on to the visitor, and that the story is a "good story", in that it respects basic cause-effect relationships and makes sense to humans. Therefore the Bayesian network described in the previous paragraph needs to be extended with additional nodes for content selection [figures 22, 23]. The additional "good story" node, encodes, as prior probabilities, the curator's preferences about how the story for each object should be told.
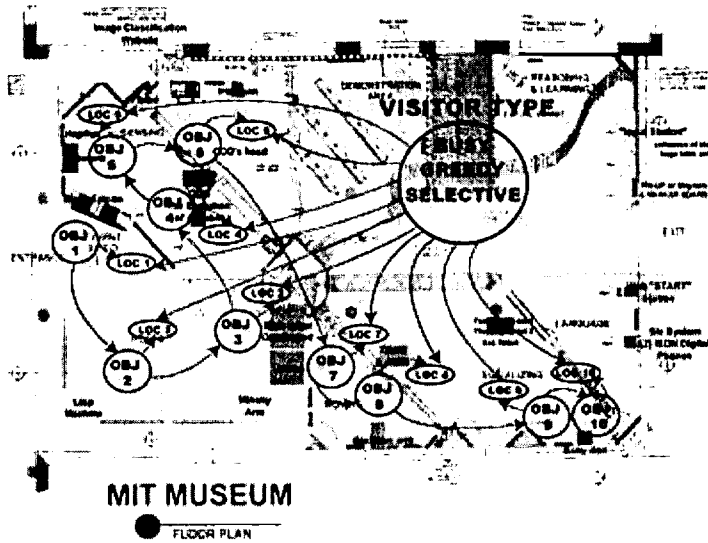
Figure 18. The geography of the exhibit needs to be reflected into the topology of the network
(click for a detailed image)

19

Figures 19, 20. Visitor wearing the museum wearable and receiving an audiovisual story about the displayed artwork (picture in picture).

Cog - Greedy Type

Cog- Busy Type

It- Greedy Type

It - Busy Type

Kismet: Selective Type

Kismet: Busy Type

Figure 21. Storyboards from various video clips shown on the museum wearable's display at MIT Museum's Robots and Beyond Exhibit

CATEGORIES / TITLES LENGTH IN SECS		Bitinside 021	Bitintro 090	Cogdrum 083	Cogfuture 043	Coghistory 051
Description	DSC	0.7	0.2	0.3	0	0
History	HST	0.1	0	0	0	0.5
Context	CTX	0	0	0	0	0
Biography	BIO	0.1	0.2	0	0	0.1
Process	PRC	0	0	0.6	0	0
Principle	PNC	0	0.4	0.1	1	0.2
Form & Function	FAF	0	0.2	0	0	0.2
Relationships	REL	0.1	0	0	0	0
Impact	IMP	0	0	0	0	0
Total P		1	1	1	1	1

CATEGORIES / TITLES LENGTH IN SECS		Cogintro 041	Dexdesign 114	Dexdintention 034	Dexdintro 072	Dexdtiffness 096
Description	DSC	0.8	0.1	0.3	0.5	0.4
History	HST	0	0	0	0	0
Context	CTX	0	0	0	0.2	0
Biography	BIO	0	0.2	0	0	0
Process	PRC	0.2	0.3	0	0	0.2
Principle	PNC	0	0	0.2	0.1	0.4
Form & Function	FAF	0	0.4	0.5	0	0
Relationships	REL	0	0	0	0	0
Impact	IMP	0	0	0	0.2	0
Total P		1	1	1	1	1

CATEGORIES / TITLES LENGTH IN SECS		Itshort 026	Klscynthiabio 067	Klsdevelop 066	Klsfacesensor 160	Klsintro 066
Description	DSC	0.8	0	0.1	0.3	0.3
History	HST	0	0	0.2	0	0
Context	CTX	0	0	0	0	0
Biography	BIO	0.2	0.9	0	0	0
Process	PRC	0	0	0.1	0.5	0.3
Principle	PNC	0	0	0.2	0	0.3
Form & Function	FAF	0	0	0.4	0.2	0.1
Relationships	REL	0	0.1	0	0	0
Impact	IMP	0	0	0	0	0
Total P		1	1	1	1	1

CATEGORIES / TITLES LENGTH IN SECS		Klssocial 183	Leg3Dbiped 055	Legblobby 021	Legflamingo 043	Legflamjerry 118
Description	DSC	0.1	0.3	1	0.5	0.4
History	HST	0	0	0	0	0
Context	CTX	0.1	0	0	0	0
Biography	BIO	0	0	0	0	0.2
Process	PRC	0.3	0.6	0	0.5	0.3
Principle	PNC	0.5	0	0	0	0.1
Form & Function	FAF	0	0	0	0	0
Relationships	REL	0	0.1	0	0	0
Impact	IMP	0	0	0	0	0
Total P		1	1	1	1	1

Table 1. Selected segments cut from the video documentation available for the MIT Museum's Robots and Beyond Exhibit. All segments have been assigned a set of probabilities which express their relevance with respect to nine relevant story themes or categories.
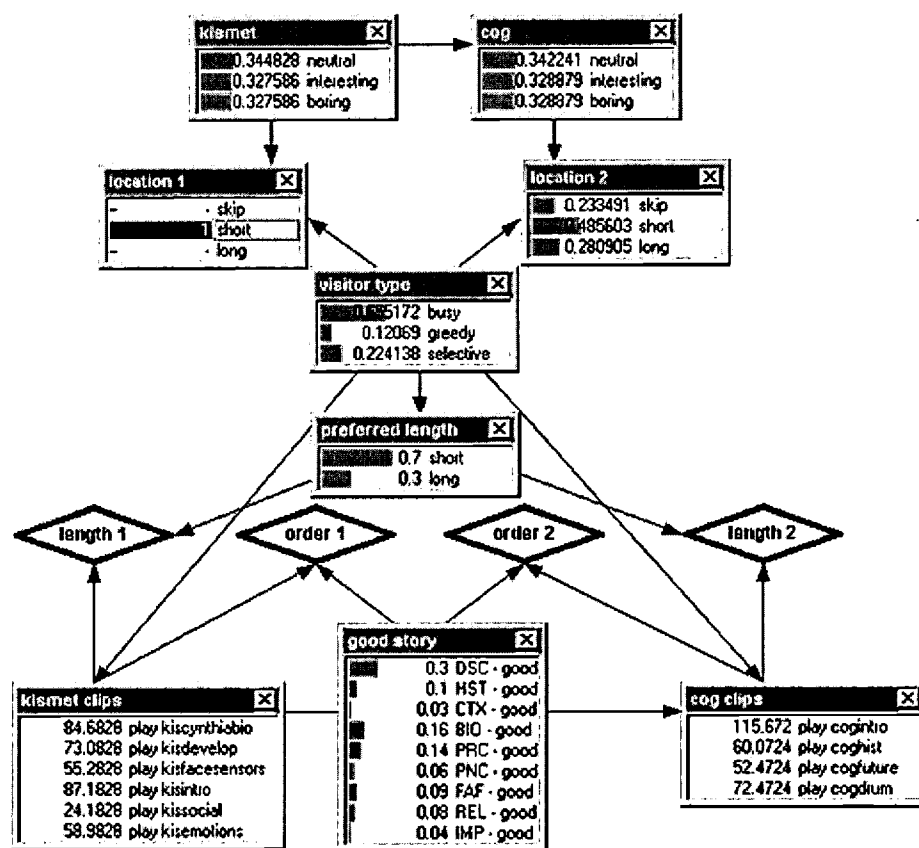
Figures 22, 23. Extension of Bayesian network to perform content

22

selection.

A study of how content is distributed geographically along the exhibit, both in two and three dimensions, was also performed. The purpose of this study was to visualize the different stories for different visitors edited by the museum wearable as paths through the hyperspace of content in the exhibit. The 2D study shows colored pie charts in the vicinity of the twelve tracked objects at the museum. Each pie chart represents the content available for the corresponding object. The size of the pie chart is proportional to the amount of content available for that object. The size of the colored slices of the chart represents the contribution of each story bin to the content available for the object [figure 24]. The 3D content map provides a visualization of how the content bins contribute to create a storyscape specific to this exhibit. It contains color coded vertical columns (a color for each content bin) whose height is proportional to the amount of content that each bin contributes to for the corresponding object [figures 25, 26].

For content personalization the system should be able to infer an interest profile of the visitors, in addition to their type as they wander along the exhibit gallery. An interest profile in the context of this research means a rating of preference for the story themes given in table 1. A GSR (Galvanic Skin Response) sensor can potentially give very useful information to the museum wearable. This sensor responds to skin conductivity and is often used in the medical and psychological field as an aid to monitor an individual's level of excitement or stress (Healey, 1999). If for example the GSR sensor measures a train of peaks when the wearable is playing a segment with biographical information about the portrayed artist, the system can infer, with a certain probability, that the visitor has a strong interest in this topic i.e. biography. It will then update the visitor interest profile with the gathered visitor preferences. The probabilistic framework offered by the Bayesian network approach is particularly relevant for this type of sensor. For example, the sensor could measure excitability for other reasons than that a compelling video segment being shown, such as meeting a friend, or recalling something that happened earlier during the day. These "false positive" data points would be modeled as "noise" intrinsic in the GSR measurements. The additional GSR sensor has been tested in a computer-based simulated environment. In this case, the decision node for content selection also needs to take into account the visitor's preferences, which compete with the curator's ordering preferences to assemble the best matching audiovisual story for each object.
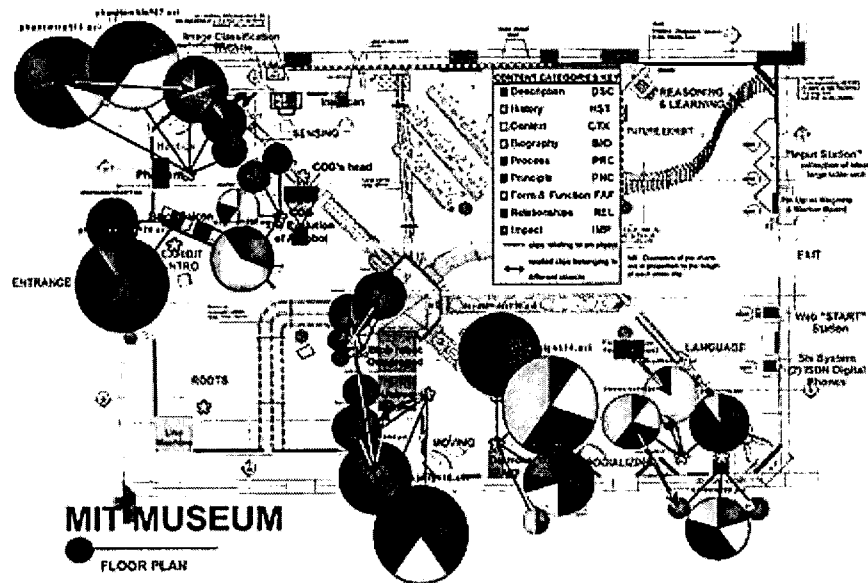
Figure 24. Two dimensional representation of content distribution for
MIT Museum's Robots and Beyond Exhibit.

BEST COPY AVAILABLE

24

Figures 25, 26. Three dimensional representation of content distribution for MIT Museum's Robots and Beyond Exhibit.

6. Potential Impact of the Museum Wearable on Exhibit Design

Potential changes and improvements that the museum wearable can produce for the space design of the MIT Museum's Robots and Beyond exhibit became obvious in the course of this project.

At the current exhibit the posters and labels occupy half of the available exhibit wall space, and while they certainly provide useful information, they require long stops for reading, take useful space away from other interesting objects which could be displayed in their stead, and are not nearly as compelling and entertaining as a human narrator (a museum guide) or a video documentary about the displayed artwork. The tracking data and our observation of museum visitors also revealed that people do not spend sufficient time to read all of what is described in the posters to absorb the corresponding information. A great deal of the space occupied by the posters and text labels is therefore wasted, as most people don't take advantage of information provided in a textual form.

The video stations, located in each section of the exhibit, complete the narration about the artwork by showing the robots in motion and by featuring interviews with their creators. While the video stations provide compelling narrative segments, they are not always located next to the object described, and therefore the visitor needs to spend some time locating the described objects in the surrounding space in order to associate the object to the corresponding narrative segment. The video stations detract attention from the actual objects on display, and are so much the center of attention in the exhibit that the displayed objects seem to be more of a decoration around the video stations than the actual exhibit.

BEST COPY AVAILABLE

The potential improvements to the exhibit layout offered by the museum wearable are summarized as follows:

1. There would be no more need to have so many posters and text labels, as the corresponding information could be provided in a more appealing audio visual form, in a video documentary style by the museum wearable. The space now made available by eliminating the large posters could be used to display more robots, which are the true protagonists of the exhibit. Typically most exhibits have to discard many interesting objects as there is not enough physical space available in the museum galleries for all objects. Therefore making more space available is a clear advantage provided to the exhibit designer and the curator. Figures 27, 28 show how the posters at the entrance of the MIT Museum's Robots and Beyond exhibit can be replaced by more objects to be seen and appreciated by the public. These images are extracted from a three dimensional Alias Wavefront's Maya 3 animation realized to visualize the wearables' potential impact on the exhibit's space [figure 29].

2. Visitors would be better informed, as the information currently provided by the posters is mostly neglected by the public. The same information would instead become part of the overall narration provided by the wearable, and it would be better absorbed and appreciated by the public.

3. The video kiosks would no longer be necessary because the same material would be presented by the museum wearable. The robots would be again the center of attention for visitors, as the wearables allows both the real world and the augmented audiovisual information to be seen at the same time as part of the wearer's real surround view. This would again make more space available for additional objects to be displayed.

The fact that the museum wearable presents audio visual material together with the corresponding object, rather than separately in space and time, and within the same field of view of the visitor, thanks to the private-eye display, is also of great importance. While no studies have been conducted yet on the quality and effectiveness of the learning experience offered by the museum wearable, there is reasonable hope that synchronous and local information provided while actually looking at the object described by the wearable can make a longer and more effective impression on the visitor. With this device curators may be able to present a larger variety of more connected material in an engaging manner within the limited physical space available for the exhibit.

Figures 27, 28. Potential impact of the museum wearable on the current exhibit layout: the posters in the Roots section (above) are replaced by new objects (below).

Figure 29. 3D model: start frame of animation of a visitor at the exhibit.

7. Related Work

Oliver (http://www.media.mit.edu/~nuria/dypers/dypers.html; Schiele, 1999) developed a wearable computer with a visual input as a visual memory aid for a variety of tasks, including medical, training, or education. This system records small chunks of video of a curator describing a work of art, and associates them with triggering objects. When the objects are seen again at a later moment, the video is played back. The museum wearable differs from the previous application in many ways. DYPERS is a personal annotation device, and as opposed to the museum wearable, it does not attempt to perform either user modeling or a more sophisticated form of content selection and authoring. It does one-to-one associations between triggering objects and recording or play back of clips. Besides general training, is used specifically in the museum context to allow a visitor to record salient moments of the explanation by a human guide to later replay them in the context of an independent visit to a museum, without a guide. The museum wearable, in contrast, focuses on estimating the visitor's type and interest profile to deliver a flexible user-tailored narrative experience from audio/video clips that have been prerecorded. These clips or animations would usually be part of the museum's digital media collection. As opposed to DYPERS, it does not have the ability to record new content to be played out at a later time. Its purpose is to create for the visitor a path-driven personalized and immersive cinematic experience which takes into account the overall trajectory of the visitor in the museum, the amount of time that visitor stays to look at and explore the objects on display, to select a personalized story for the visitor, out of several possible digital stories that can be narrated.

Höllerer and Feiner (1997) have built a university campus information system, worn as a wearable computer. This device is endowed with a variety of sensors for head tracking and image registration. Both the size of the

28

wearable, mounted on a large and heavy backpack, as well as the size of the display, are inappropriate use in a museum visit.

Various groups are working to augment the museum visit with mobile devices that are not wearable computers but handhelds, which do not have a private-eye or head-mounted display but rely on the handheld's screen for visual communication. One of the main drawbacks of such devices is that the visitor is obliged to toggle his/her attention between the objects on display and the handheld's screen, alternatively looking ahead towards the objects and then down to the screen. The private eye of the museum wearable instead allows the visitor to have a true augmented reality experience by presenting the viewer with a fused image which mixes together the real world and the computer augmentation either as a picture-in-picture effect or as two superimposed layers of information.

Amongst the handheld based projects, Spaspjevic and Kindberg (2001) describe the electronic guidebook, currently under development at the San Francisco Exploratorium Science Museum. This device uses a combination of infrared and RFID location sensors to give visitors the ability to either view or bookmark Web pages which provide additional description information on the objects of the exhibit. This portable device is used mainly to record the visitors' path through the exhibit, typically as a group, so that later the visit can be discussed and commented upon in a classroom setting.

The European HIPS project proposes a user-centered approach to information delivery in museums (Broadbent and Marti, 1997). As participants to HIPS, Oppermann and Specht (1999) describe an adaptive system which requires the user's intervention to personalize the presentation. This however can be quite disruptive of the visitor's museum experience, in a way similar to those multi-path DVD movies which stop at all turning key point and ask the viewer to choose an option before continuing. A system such as the museum wearable, capable of inferring the user's preferences with a mathematical model able to use the sensors' information as cues and the curators' knowledge of their public as a guideline, , can be indeed less disruptive.

The MUSEpad project (Kirk, 2001) is developing an original mobile device to enable visitors with disabilities to customize and optimize their learning and leisure experiences in museums. Their feasibility study examines user personalization to allow for adaptation to the needs of different users, whether by providing video-only, audio-only or magnified images for visitors with poor vision.

Aoki and Woodruff (2000) describe an electronic guidebook prototype that facilitates social interaction during the museum visit. Using a touch sensitive screen on the handheld and a photograph-based interface that the visitors can click on, the handheld offers a text or audio description of the selected object, according to the user's preferences. Visitors are then provided a mechanism to hear each other's audio selection. Such research on mobile augmentation devices which can be used collectively as a group is quite important for the future of this field.

The author showed an early prototype of the museum wearable called "Wearable City of News", as a demonstration for the SIGGRAPH 99 Millennium Motel, where hundreds of users tested the system for seven full days (http://www.siggraph.org/s99/conference/etech/projects.html). The system featured a jacket with an embedded computer and head-mounted display to show visitors Web pages and video previews of the Millennium

29

Motel demonstrations using a short range but low power infrared location system [figures 30, 31].

Figures 30, 31. Early version of the museum wearable at the SIGGRAPH 99 Millennium Motel.

8. Discussion and Future Work

The museum wearable fuses the audiovisual documentary, which illustrates and extends an exhibit, with the visitor's path inside that exhibit, using a wearable computer. By having the public use this device, curators and exhibit designers can accomplish multiple goals simultaneously: they can have objects narrate their own story; they do not needs special rooms to show audiovisual explanations about the exhibit as with the wearable the narrative is unfolded by the visitor's path in the museum; they can show more artwork than what is physically on display, including video, images, audio, and text about other important objects for the exhibit (these usually do not see the light because of the physical limitations of the available space); they do not need to disseminate panels with textual explanation or video monitors along the aisles of the exhibit as that information can now be tailored to each individual visitor; they can personalize the audiovisual explanations provided to the public based on the visitor's type and exploration strategy.

The museum wearable provides more than a simple associative coupling between inputs and outputs. The sensor inputs, coming from the long range indoors infrared positioning system, are coupled to digital media outputs via a user model, and estimated probabilistically by a Bayesian network. The ability to coordinate and present the visual material as a function of the visitor's estimated type (i.e. busy, greedy, or selective types, or other appropriate types), seamlessly, appropriately, and in conjunction with the path of the wearer inside the exhibit, is an important feature of this device. Bayesian networks have the additional advantage that they allow to us encapsulate our human knowledge about the context of use of the museum wearable (a particular exhibit, a trade show) in appropriate nodes of the network.

With respect to the traditional museum audio tour, the museum wearable introduces the following innovations: it does not constrain the visitor to follow any predefined sequential path in the museum, or to continuously press buttons, but it relies in its sensing system to find the visitor's location and respond; consequently, it adds a layer of visual augmentation, not just auditory; through Bayesian network based user modeling it provides personalized content to each visitor, as a function of the estimated visitor type.

An experimentation phase at the museum should follow the current research. A procedure must be set to establish if and how the museum wearable does actually enhance learning and entertainment at an exhibit, or how the content shown does actually match the visitor's preferences.

The research here presented can also be expanded in various ways. One direction of work is to find ways to get to know the visitors better so as to target content presentation more accurately towards their level of knowledge or competence. Asking visitors to fill out lengthy questionnaires upon entering the exhibit may not be practical. It is instead desirable to experiment with additional sensors, such as the GSR, a motion sensor (accelerometer), or a tiny video camera to obtain a more accurate estimate of the visitors' interest profiles and levels of attention. More visitor tracking data could be gathered at the museum site, to eventually infer more visitor types than the ones described in this document, and compare them with the more sophisticated visitor typologies discussed in the museum literature.

The museum wearable can also become a very useful tool to gather visitor tracking data in the museum. For example, rather than coming up with a set of visitor types from the museum literature, one could adopt the opposite approach of "inferring" the visitor types from a statistical analysis of the tracking data (path and stop duration) gathered by visitors with the museum wearable. This information would help the curator, exhibit designer, and the modeler of the interactive museum experience to refine their knowledge about visitor types for a specific exhibit. Similarly, by analyzing the posterior values of the object nodes, the curator and the exhibit designer could see which objects are the most interesting or boring for the visitors, and change the exhibit layout accordingly.

An important extension to the museum wearable would allow it to support visitors who want to come to the museum as a group and have the freedom to comment and discuss the artwork amongst themselves. A simple modification to the current prototype would be to add a small microphone capable of detecting when the visitor is talking, and then automatically pause the narration. If for example the group of visitors is composed of high school students, it would be useful, for learning purposes, to make each visitor