International Cultural Heritage Informatics Meeting: 2007

October 24-26, 2007
Toronto, Ontario, Canada

Paper: Situating Approaches to Museum Guides for Families and Groups

Ron Wakkary, **Kevin Muise**, **Karen Tanenbaum**, **and Marek Hatala**, Simon Fraser University and **Leora Kornfeld**, Ubiquity Interactive, Canada

Abstract

In this paper, we analyze the current state of museum guide technologies and applications in order to develop an analytical foundation for our future research in an adaptive museum guide for families. We have focused our analysis on three critical areas of interest in considering group and social interaction in museums: tangibility, the role of tangible user interfaces; interaction, visit types and visit flows; adaptivity, user modeling approaches.

Keywords: museum guides, adaptive museum guides, tangible user interfaces, user models, group interaction, mobile, family

Introduction

Museum technologies have increasingly become the focus of research in such areas as ubiquitous computing, tangible computing and user modeling. Since the advent of audio-based museum guides, much research and development has been placed on increasing the technological capacity of augmenting the museum visit experience. Early prototypes such as Bederson (Bederson 1995) provided evidence that it is possible to support visitor-driven interaction through wireless communication, thus allowing visitors to explore the museum environment at their own pace. Current prototypes and fully functional systems are much more complex, supporting a variety of media, adaptive models, and interaction modalities. However, as Bell notes in her article on museums as cultural ecologies; "The challenge here is to design information technologies that help make new connections for museum visitors" (Bell 2002). Bell further argues for the importance of *sociality in* museums where visits are as equally social and entertaining as they are educational.

At this stage in our research in adaptive museum guides we are exploring how best to address issues of social engagement and play and learning for groups of visitors such as families. Families are by far the most common visitor type to science, history, and natural history museums. In this paper, we describe our theoretical analysis of the current state of museum technologies and adaptive museum applications in order to provide a detailed understanding in support of our current research goals. We've narrowed our focus to three critical trends in adaptive museum guides that we believe will form the theoretical anchors of our future research. The areas include: tangibility, interactivity, and adaptivity. We discuss these trends in order and follow with a discussion on the impacts on our research. For the reader, we believe this paper provides a critical summary of electronic museum guides of the last decade and outlines key areas of concerns for academic researchers.

Situating Tangibility In Museums

Tangible user interfaces (TUI) imbue physical artifacts with computational abilities. Like no other user interface concept it is inherently playful, imaginative and even poetic. In our previous research in adaptive museum guides we developed a prototype system known as ec(h)o. In this project, the simple physical display devices and wooden puzzles at the natural history museum where we conducted ethnography sessions inspired us. As a result, the user interface for ec(h)o was a TUI that coupled a wooden cube with digital navigation and information. In ec(h)o, museum visitors held a light wooden cube and were immersed in a soundscape of natural sounds of and information on the artifacts on display (figure 1). Visitors navigated the audio options presented to them by rotating the cube in their palm in a direction that corresponded to the spatial location of the audio they were hearing (we will return to this project below but for a detailed account see (Wakkary and Hatala 2007)).
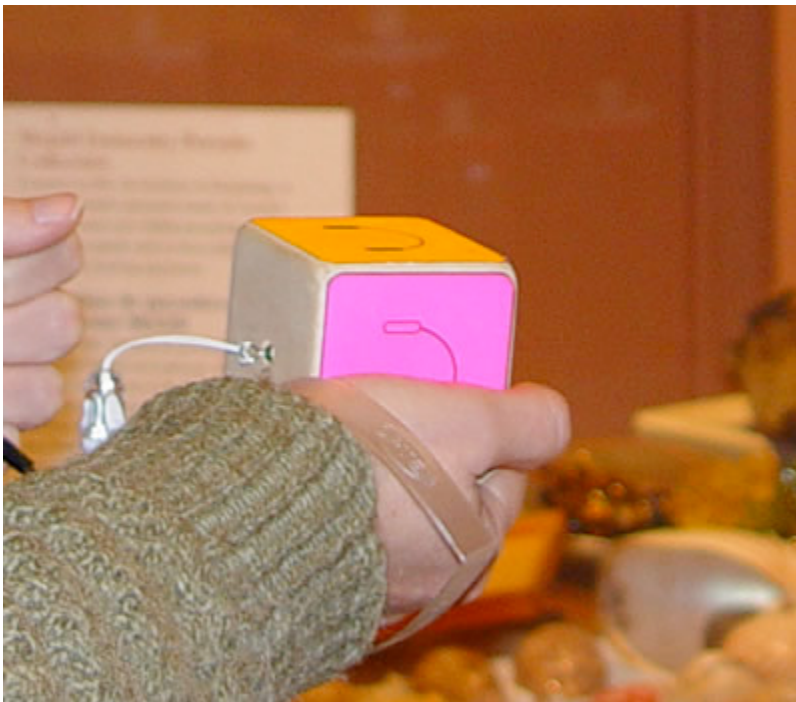

Fig 1:the tangible-user interface of ec(h)o

In our experience with TUIs, they bridge between the virtuality of the museum guide system and the physical surroundings of the exhibition. As such, the adaptive museum guide firstly becomes more integral to the physical ecology of the exhibition including artifacts, display systems, and architecture rather than being a separate technology. It is important to note that our social interactions are in large part mediated by objects and environments as much as by direct contact with others (Kaptelinin, Nardi et al. 2006). Awareness of context is critical to sociality. Secondly, with TUIs our understanding of interacting with technology is informed by our rich experience with physical artifacts and surroundings – since we can leverage existing embodied and cognitive understanding. Analytically, Kenneth P. Fishkin (Fishkin 2004) provides a taxonomy that allowed us to explore these two factors further.

Fishkin's taxonomy is a two-dimensional space across the axes of *embodiment and metaphor*. Embodiment characterizes the degree to which "the state of computation" is perceived to be in or near the tangible object. As we discussed, it expresses the integration of computation with the physical artifact and environment, Fishkin details four levels of embodiment:

1. *distant* – representing the computer effect is distant to the tangible object;
2. *environmental* – representing the computer effect is in the environment surrounding the user;

3. *nearby* – representing the computer effect as being proximate to the object;
4. *full* – representing the computer effect is within the object.

Along the second axis, Fishkin uses metaphor to depict the degree to which the system response to user's action is analogous to the real-world response of similar actions – the existing embodied and cognitive mappings. Fishkin divides metaphor into *noun metaphors* , referring to the shape of the object, and *verb metaphors*, referring to the motion of an object. Metaphor has five levels:

1. *none* – representing an abstract relation between the device and response;
2. *noun* – representing morphological likeness to a real-world response;
3. *verb* – representing an analogous action to a real-world response;
4. *noun + verb* – representing the combination of the two previous levels;
5. *full* – representing an intrinsic connection between real-world response and the object which requires no metaphorical relationship.

In figure 2, we have applied Fishkin's taxonomy to ec(h)o. *Embodiment would* be considered "environmental" since the computational state would be perceived as surrounding the visitor given the spatialized audio display output. In regard to *metaphor*, the ec(h)o TUI would be a "noun and verb" since the wooden cube is reminiscent of the wooden building blocks and the motion of the cube determines the spatiality of the audio, as turning left, in the real-world, would allow the person to hear on the left. If we consider the visitor's movement, the embodiment factor would still be environmental and we'd have to consider the visitor's body as being "full" in Fishkin's use of metaphor. And so in representing the entire system, we plotted ec(h)o between "noun + verb" and "full" on the metaphor axis.
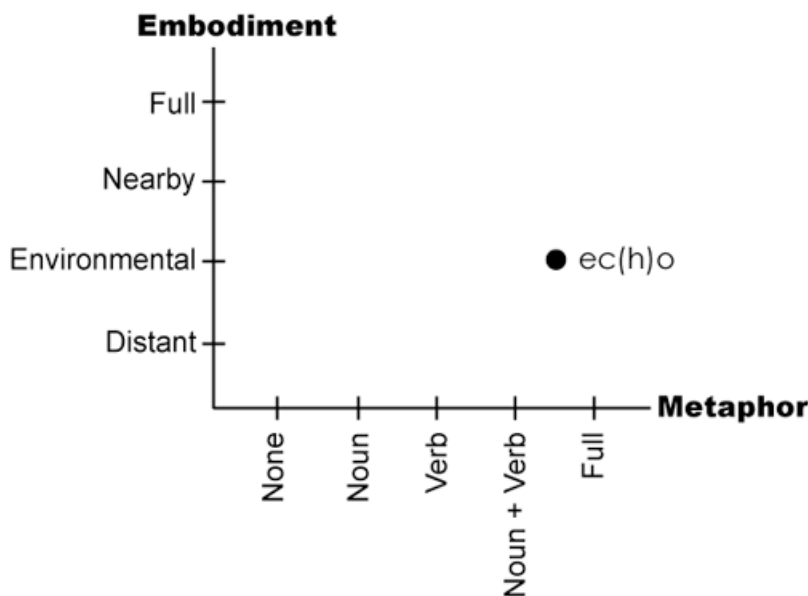

Fig 2: ec(h)o plotted in Fishkin's TUI taxonomy

It is natural to think that optimal TUIs would be "full" in both *embodiment and metaphor dimensions*, and Fishkin suggests as much. Yet we caution that in the case of museum guides, and in particular, when considering the sociality of museum visits, *embodiment is* between people, technology and the environment. Accordingly, TUIs in museums are optimal when computation is perceived to be *environmental and nearby rather* than solely within the device.

In summary, we believe tangibility is critical in considering the social dimensions of museum visits in museum guides. TUIs can integrate the computational and physical affordances of the museum visit experience, and visitors can leverage embodied and cognitive models of interaction in incorporating new technology. Visitors benefit from augmentation of computation yet maintain awareness and embodied connection to their

surroundings, which ultimately supports social interaction.

Situating Interactivity In Museums

Interactivity can be understood in many ways. We focused on individual or group experiences in museum visits and models of content delivery. In this section, we review a range of systems from the last decade with a particular eye on how these past approaches provide insights into designing for families in which sociality and group activity are paramount. In reviewing the various systems we have developed a matrix that compares the visit type (individual/group) to visit flow (information delivery/game-interaction) which aims to uncover the major factors that affected the design of the various systems.

In visit types, we refer to *individual as* systems that target interaction with single individuals, whereas *group refers* to interaction models aimed at groups of individuals. *Information delivery refers* to the approach of an information corpus or repository that is presented to the visitor whether adaptively or by user selection. Providing the visitor with the ability to interact with the exhibit content in a playful manner we've referred to as *game-interaction*. Often the *game-interaction approach* will use games as a means to educate the visitor, as opposed to providing information about an artifact. By employing a cross matrix of both types of categorizations, it becomes possible to understand how both past and current museum guides mediate the museum experience.

Many of the systems we reviewed were *information delivery in* style, while also falling under the *individual category*. Typically they involved a Personal Digital Assistant (PDA) and an audio/visual interface such as the Blanton iTour (Manning and Sims 2004). Among the first of such guides was the HyperAudio project (Petrelli and Not 2005), which used an adaptive model that we will discuss again in the next section. In the HyperAudio system, individual visitors were encouraged to walk about the exhibition with the handheld device and headphones, stopping at various exhibits to learn more about specific artifacts in their surrounding. As the authors, Petrelli and Not note,

the presentation (audio message and hypermedia page) would be adapted to each individual user, taking into account not only their interaction with the system, but also the broad interaction context, including the physical space, the visit so far, the interaction history, and the presented narration (Petrelli and Not 2005).

The presentation that is displayed also provides a link that the participant can click on with a palmtop pen to gather further information about the artifact of interest. By doing so, the visitor is provided a map of the museum on which the location of the new artifact is displayed, to allow the visitor to see the artifact in person if he desires. The authors also describe their interest in keeping the graphical user interface to a minimum so as not to distract the visitor. It is for this reason that audio is chosen as the main channel of information delivery.

The PEACH project (Stock, Zancanaro et al. 2007) provides the visitor with a digital character on their PDA that provides information on various artifacts within the exhibition. The PEACH guide also contains rich media in the form of video close-ups of frescoes and detailed descriptions of paintings. Unique to this system is the availability of a printout, which provides the visitor with an overview of the exhibits he encountered while at the museum.

We've observed in a commercial system developed by a member of our team as part of Ubiquity Interactive, known as VUEguide that rich media images can at times become more engaging than the authentic object on display. We believe the virtual image can be easily integrated into a narrative world of information delivery. For example, interactive images of Bill Reid's 'Raven and the First Men' sculpture engaged visitors deeply with the screen and encouraged a 'back and forth' exploration between the actual and the virtual (see figure 3).

Fig. 3 interactive image of a Haida sculpture that visitors found engaging
in the VUEguide system

Berkovich et al (Berkovich, Date et al. 2003) developed the Discovery Point prototype, a four-button device that delivers audio to the museum visitor. The device allows the visitor to listen to a short audio clips about artifacts that are controlled by the visitor. The visitor can also create a virtual souvenir by pressing the "MailHome" button, which adds the artifact in question to their personal Web site created upon entering the museum. The device functions without the use of headphones. Instead, audio is delivered through specialized speakers which direct audio to a precise area around the artifact, so as not to disturb other museum visitors.

The Sotto Voce system, produced by Aoki et al (Aoki, Grinter et al. 2002), provided the first instance where researchers focused on group activity. The system contained an audio sharing application called *eavesdropping that* allowed paired visitors to share audio information with each other while on an information delivery tour. In designing the application with three entities in mind (the information source, the visitor's companion and the museum space), the authors remarked that the system showed that the visitors used the system in creative ways, and with social purpose, while presenting the opportunity for co-present interaction.

In recent years, research has continued on group-based museum tours as an area of interest, but orientation has shifted from information delivery tours to *game-interaction activities*. This is evident in the CoCicero project implemented in the Marble Museum in Carrera, Italy (Laurillau and Paternò 2004). The CoCicero prototype focused on four types of group activities; (i) *shared listening* – similar to the Sotto Voce system, (ii) *independent use* – to allow individuals to choose not to engage in group-based tours, (iii) *following* – to allow an individual to lead other members of a group, and (iv) *checking in* – which allows members in a group to know how others are doing through voice communication while not being physically present. The authors state that communication; localization, orientation and mutual observation are four elements that are key to a collaborative visit. The guide functions by providing museum groups a series of games, such as a puzzle and multiple choice questions, which require the visitors to gather clues through viewing the exhibits within the museum.

Similarly, the ARCHIE Project (Loon, Gabriëls et al. 2007) has developed a learning game for school children that allows visitors to trade museum-specific information to gain points in order to win a game. With each player having a specific role that he plays, the visitor must understand various levels of information gained from exploring the museum in order to improve his score. Though the ARCHIE project is only in its prototype stage, it is clear from their initial findings that the *game-interaction approach* does foster user interest in museum content. Our own system, ec(h)o, is unique in relation to other systems. It is exclusively an individual

type system (which we now see as a significant drawback) yet it used a *game-interaction approach*. Our interface aimed for an interaction based on an open-ended game qualities or what we referred to as *play*. *Play took* on two forms: (1) *content play* in the delivery of information in the form of puns and riddles; (2) *physical play* that consisted of holding, touching and moving through a space; simple playful action along the lines of toying with a wooden cube.
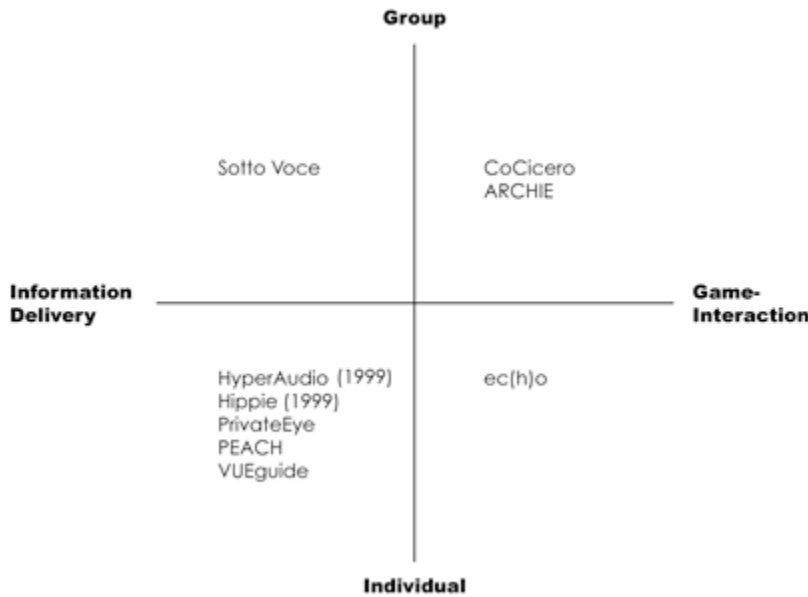

Fig 4: Cross matrix of visit flow and visit type

In summary, it is clear that the majority of research has been in the area of *information delivery* and *individual* visit type. See figure 4 where we plotted the different systems we reviewed. Sotto Voce is among the first to explore co-visiting or group interaction with a museum guide within an information delivery context. Group interaction is a new area and chronologically represents a trend, especially when combined with a *game-interaction* approach.

Situating Adaptivity In Museums

This section surveys those projects that deployed user modeling. User modeling is the use of artificial intelligence software techniques to construct conceptual models for users that enable predictions and responsiveness to user interactions. The projects we review include: Hippie, the Museum Wearable, ec(h)o, and HyperAudio. We have also analyzed a system that is currently Web-based, known as CHIP. Each system offers the museum visitor some form of personalized content such as audio or video clips, text, and images, and each uses aspects of visitor interaction with the system to achieve customization. The details of each project are discussed first, where we provided a detailed account of each project's approach to adaptivity. The overview will be followed by an analytical discussion that includes a cross matrix used to frame the current research in the field.

Adaptive Systems

Earlier we introduced the HyperAudio system from 1999 (Petrelli and Not 2005). Audio content generation was based on a *visit model* that looked at the group each user was in, whether it was a first or repeat visit, how long they planned to stay and what kind of interaction type they preferred. Underlying the model was a simple boolean variable indicating interest in each item, based on whether the user remained in front of an object after the relevant presentation finished, or whether they turned the presentation off.

Another early attempt to develop a context-sensitive, adaptive museum guide, was the Hippie project

(Oppermann and Specht 1999) and (Oppermann and Specht 2000). Hippie contained models for 3 distinct components: 1) a static domain model (objects to be presented and processed about), 2) a static space model (physical space), and 3) a dynamic user model. In constructing the user model, the authors assumed that visitors had a "stable interest trait structure", but that environmental and contextual factors played a role in the actual activation of that structure at any given moment. The goal of their user model was "to predict the information needs of a user in a given episode of a visit," thus the model made inferences regarding the next exhibit to visit and the next piece of information to present. Interest in particular exhibits was inferred by time spent there as well as the number of information items selected. They also suggest using "differentiated navigation behavior" to indicate interest, i.e. viewing artworks from multiple locations rather than one indicates more interest. Guide content was classified according to type, and used to infer what kind of information the visitor was interested in. The HIPS project also applied Veron & Levasseur's 1999 ethnography work in museums, which identified four physical movement patterns of visitors: ant, butterfly, fish, grasshopper (Marti, Rizzo et al. 1999).

- Ants proceed methodically through the entire museum space, looking at everything;
- Butterflies fly around from work to work based on interest level;
- Fish swim through the space quickly, glancing at things in a cursory fashion;
- Grasshoppers also bounce around, but with less of a defined idea of what they are interested in.

Visitor type and information preference was proposed as a way to assemble appropriate length audio clips for each individual.

While the Hippie system asked the users what areas they wish to explore, other approaches did not give this level of control to the user, such as the Museum Wearable, which assembles and delivers personalized content to the user without explicit user interaction (Sparacino 2002a) and (Sparacino 2002b). A Bayesian network is used for user classification & decisions about content assembly. Sparacino uses a ternary classification of visitors as busy (cruise through everything quickly), selective (skip some things, spend long on others), greedy (spend a long time on everything). Data was collected and used to set prior probabilities of a model based on path and pause duration patterns. Upon approaching a new exhibit, a sequence of content clips was dynamically assembled using user type to determine how long it should be.

Th ec(h)o project provided visitors with a choice of audio clips throughout their museum visit (Hatala and Wakkary 2005). The ec(h)o user model captured two basic kinds of information about the museum visitor using the guide; interaction history and user behaviour information. Interaction history recorded movement through the museum space along with the sound objects selected for listening, whereas user behaviour information kept track of the user's interests. Interest was set by the user's explicit choice from a number of different concepts at the beginning of the museum tour and then updated based on the exhibits they paused at during their exploration of the museum space and the choices they made regarding which audio content to listen to. If the user showed special interest in a specific concept or two, those concepts would become highly weighted in the model and content related to them would be offered frequently. However, to avoid the complete stereotyping of a visitor, a "variety" element was also included in the recommendation algorithm so there was the possibility of sound objects being offered that would allow the user to move away from their heavily weighted interests.

The most recent project to use adaptability in museums is the CHIP project, which is currently Web-based only. This project allows a user to generate a personal profile via an online rating system for artwork (Aroyo 2007). Images in the system are semantically annotated for artist, period, style, visual content and themes. As users rate them with 1-5 stars, the user model is refined to better reflect their preferences. The classification ratings (e.g. 4 stars for Impressionism) are inferred from the explicit work ratings, and can be viewed at any time. If the user disagrees with the system's calculation of the classification rating, they can adjust it manually. The user can also ask to see "recommended works" which pulls out unrated artwork that the model believes the user would rate highly (4-5 stars). Users can query as to why the system recommended this work, and will be provided with the individual classification ratings that underlie the suggestion. The next step in the project is to integrate this with technology in the physical museum, so that users can access their profile in the exhibit.

The PEACH project also allows the user to actively participate in the construction of her user-model through a

widget (Stock, Zancanaro et al. 2007). The "like-o-meter" widget allows the visitor of the system to state whether she likes or dislikes a given museum artifact, which affects the amount of information the system provides for subsequent related artifacts. The authors suggest that their widget was clearly understood and enjoyed by visitors.

Analysis

The goal of this analysis is to get a sense of what is common practice in user modeling, what is more uncommon and experimental, and what has not been attempted yet. One of the first aspects to consider is what kind of data the systems gather and use as input to the model. There are two basic categories of input commonly used: physical data and content-based data. Physical data includes things like knowledge of the visitor's current location (Hippie, Museum Wearable, ec(h)o, and HyperAudio), of their overall path through the museum space (Hippie, Museum Wearable), and of their stop duration at each specific exhibit (Hippie, Museum Wearable, HyperAudio). Content-based information includes knowledge of what content the visitors have listened to or selected (Hippie, Museum Wearable, ec(h)o, CHIP, HyperAudio) and how they rate that content (CHIP, PEACH). From this basic input, all of the systems extrapolate and make inferences about categories the user falls into or characteristics the user might have, based on the observed behavior. Every system infers user interest in an artifact/exhibit, usually based on their movement towards it, presence near it, or selection of it in some way. The systems all have the capacity to detect user interest in broader themes or topics, as signaled by the selection of multiple items, which have been annotated similarly. Some of the systems also include the ability to detect interest in a specific information type, as signaled by selection of specific kinds of content (the Museum Wearable, Hippie, and HyperAudio). For example, a user might repeatedly ask for artist biography content, or for details on how an artifact was constructed. The final type of inference found in these systems is the classification of visitors as a certain "user type". This includes the fish/ant/butterfly/grasshopper path classification from the Hippie project and the greedy/busy/selective user type from the Museum Wearable, both of which used movement patterns to sort users.

Content and Interaction type matrix

One important distinction amongst the ways the projects handle system input is the degree of control that the user has over the cues that the system is picking up on. Another way to think about this would be in terms of the opacity level present in the interaction. With some of the systems, it is easy for users to tell what kinds of information the system is picking up on, such as their presence next to an object or their deliberate selections within the system; we describe this approach as *transparent interaction*. Other systems have much more opaque interactions that are being picked up on, such as the path through the space or their patterns of stop duration and movement. In these cases, the user of the system may never guess that such elements are being used as input for the user model, and thus it is beyond their control to affect what the system is doing in response to them; we describe this approach as *opaque interaction*. It's possible that this opacity will yield a more natural, intuitive interaction with the system, but it's equally possible that it could result in a frustrating experience where the user does not understand why the system is responding the way it is.

With the collected and inferred data as input, the next concern is what the system presents back to the visitor as output. Several of the guides offer a set of further audio and/or video content that the user can select from (Hippie, ec(h)o). The Museum Wearable also presents audio/video content, but does not allow the user to select what they view; instead the model automatically assembles a tailored presentation for the user. HyperAudio combines these approaches by assembling a tailor made audio presentation and also allowing for individual exploration of topics via the handheld device. A third option for output is to generate recommendations of other pieces the visitor might like based on the model's understanding of the visitor (CHIP). There is a basic distinction here between systems that have *static content* and those that have *personalized content*. The Museum Wearable and HyperAudio actually tailor the content itself to the individual user. In the other systems, the content remains the same, but the order in which it is presented or the options available at any moment are tailored to the individual user.

Static Content

Hippie   ec(h)o
                     CHIP
                     PEACH

Opaque Interaction   **Transparent Interaction**

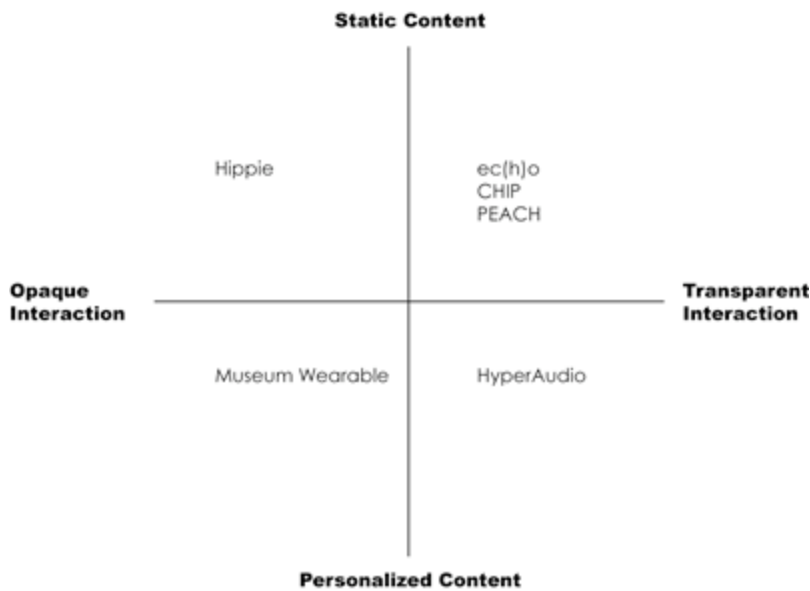Museum Wearable   HyperAudio

Personalized Content

Fig 5: Cross matrix of interaction and content types

In summary, Figure 5 shows a cross-matrix of this *static-personalized content* dimension with the *opaque-transparent interaction* dimension. As can be seen from the diagram, the majority of research thus far with user modeling in museum applications has involved static content and transparent interaction. There have been a couple of different ways in which the input/inference data affects the output. All the guides have the simple correlation of high user interest leading to more content related to that user interest. Hippie and the Museum Wearable also use user type classifications to affect the duration of the content offered to the user. Room still exists for a range of creative thinking in this area, especially with regards to how collected data can influence the system interaction and output.

Discussion

Through gathering research in this study, it is possible to understand the variety of approaches taken to augment the museum space. When faced with group-based activities, the PDA may further distract the visitor from his/her companions. The group-based applications discussed in this study revealed a shift toward a *game-interaction* approach – while all used PDA-like devices to guide the visitors through the museum, we believe TUIs may be even more effective in this regard. Research within group-based tour applications is fairly limited, and has become the focus of two newer studies (CoCicero and ARCHIE). The *game-interaction* approach seems not only to affect the way in which individuals browse through the museum, but may also affect the manner in which learning occurs. The *game-interaction* approaches discussed here encourage the visitors to find artifacts, which match specific criteria in a quest-like fashion. This type of activity might be useful in teaching visitors what singular artifacts are, but may inhibit the communication of the contextual significance that often surrounds artifacts in museums. The earliest of systems to adopt a group-based approach, Sotto Voce, note that co-present interactions should be supported. The CoCicero project takes the concept further by introducing other group-based behaviours to be supported, such as *following* and *checking-in*. As our focus is on group-based interaction, using the findings of such projects may prove helpful in designing an application that fosters group visits by families to museums.

With regards to a user modeling approach, the majority of the research lies within supporting *transparent interaction*, and *static content*, though it is difficult to state whether one method is superior to another at this time. It could be argued that a push towards even more transparent input is desirable. The CHIP project is the most *transparent* of all the interactions, allowing the visitor to adjust the user model manually and get feedback from the system as to why certain recommendations have been made for them. No such hybrid adaptive and adaptable system has been implemented within the user space itself. Alternatively, it could be fruitful to

explore the area on the extreme other side of the *transparent-opaque interaction* continuum, by creating a system where users are unable or unlikely to guess what aspects of their behavior are being used within the system. If done properly, this could yield a very immersive, natural feeling interaction. Projects that have attempted to classify visiting patterns have done so within the individual-based tour context, and there is a lack of research on museum visitor patterns for groups – a potential area of interest for us to address.

There is also little research on adapting to the group conditions of the visitor, e.g. taking into consideration the fact that person A is a mother visiting with three children while person B is a senior citizen visiting with friends. Modeling users on the level of group dynamics is a new and growing area within user modeling, one that we have explored (Wakkary, Hatala et al. 2005) and hope to bring to the museum domain. In terms of content, there is definitely space to explore in the realm of adaptive content, taking the personalization of the museum beyond simple recommendations and on-demand information access and moving it into the realm of individualized presentations.

We began the paper with a discussion on tangibility as a mediating bridge between computation and context that militates the lack of support for social interaction. Tangibility may preclude some of the directions in user modeling such as *transparent interaction* of systems like CHIP, which is Web-based after all. However, tangibility could be seen as an enabler in *game-interaction* approachesand *group* interaction, given its relative playfulness in comparison to graphical interfaces and the ease of sharing tangible devices and the affordance of leveraging past patterns of social interactions mediated by objects.

Conclusion

As shown in this paper, there are currently several approaches for adaptive museum guides under exploration. Through analyzing current contributions through the three key areas of interest, tangibility, interactivity, adaptivity, we were able to develop a comprehensive outline that provides a technical grounding for our future research in a family and social based museum guide. The majority of the current literature focuses on the interactions of a single visitor but through analyzing trends in the museum guide research, the focus appears to be shifting.

References

Aoki, P. M., R. E. Grinter, et al. (2002). Sotto voce: Exploring the interplay of conversation and mobile audio spaces. Proceedings of the ACM conference on human factors in computing systems (CHI '02), Minneapolis, Minnesota, ACM Publications.

Aroyo, L. (2007). Personalized museum experience: The Rijksmuseum use case. in Museums and the Web 2007: Proceedings, J. Trant and D. Bearman eds., Toronto: Archives and Museum Informatics. Available http://www.archimuse.com/mw2007/papers/aroyo/aroyo.html

Bederson, B. (1995). Audio augmented reality: A prototype automated tour guide. Conference companion of the ACM conference on human factors in computing systems(CHI '95), Denver, Colorado, ACM Press.

Bell, G. (2002). Making sense of museums: The museum as 'cultural ecology', Intel Labs: 1-17.

Berkovich, M., J. Date, et al. (2003). Discovery point: Enhancing the museum experience with technology. CHI '03 extended abstracts on Human factors in computing systems, Ft. Lauderdale, Florida, USA, ACM Press.

Fishkin, K. P. (2004). "A taxonomy for and analysis of tangible interfaces." Personal Ubiquitous Computing 8(5): 347-358.

Hatala, M. and R. Wakkary (2005). "Ontology-based user modeling in an augmented audio reality system for museums." User Modeling and User-Adapted Interaction, The Journal of Personalization Research 15(3-4):

339-380.

Kaptelinin, V., B. A. Nardi, et al. (2006). Acting with technology activity theory and interaction design. Cambridge, Mass., MIT Press.

Laurillau, Y. and F. Paternò (2004). Cocicero: Un système interactif pour la visite collaborative de musée sur support mobile. Proceedings of the 16th conference on Association Francophone d'Interaction Homme-Machine, Namur, Belgium, ACM Press.

Loon, H. V., K. Gabriëls, et al. (2007). Supporting social interaction: A collaborative trading game on PDA. Museums and the Web 2007: Proceedings, J. Trant and D. Bearman eds., Toronto: Archives & Museum Informatics. Available http://www.archimuse.com/mw2007/papers/vanLoon/vanLoon.html

Manning, A. and G. Sims (2004). The Blanton itour - an interactive handheld museum guide experiment. Museums and the Web 2004: Proceedings, D. Bearman and J. Trant, eds. Archives & Museum Informatics. Available http://www.archimuse.com/mw2004/papers/manning/manning.html

Marti, P., A. Rizzo, et al. (1999). Adapting the museum: A non-intrusive user modeling approach. 7th International Conference on User Modeling, UM99, Banff, AB.

Oppermann, R. and M. Specht (2000). A context-sensitive nomadic exhibition guide. Handheld Ubiquitous Computing 2000, Bristol, UK.

Oppermann, R. and M. Specht (1999). A nomadic information system for adaptive exhibition guidance. Cultural Heritage Informatics 1999: selected papers from ichim99, D. Bearman and J. Trant Eds. Pittsburgh: Archives & Museum Informatics. Available http://www.archimuse.com/publishing/ichim99/oppermann.pdf

Petrelli, D. and E. Not (2005). "User-centered design of flexible hypermedia for a mobile guide: Reflections on the hyperaudio experience." User Modeling and User-Adapted Interaction 15: 303-338.

Sparacino, F. (2002). The museum wearable: Real-time sensor-driven understanding of visitors' interests for personalized visually-augmented museum experiences. Museums and the Web 2002: Proceedings, D. Bearman and J. Trant, Eds, Pittsburgh: Archives & Museum Informatics. Available http://www.archimuse.com/mw2002/papers/sparacino/sparacino.html

Sparacino, F. (2002). Sto(ry)chastics: A bayesian network architecture for user modeling and computational storytelling for interactive spaces. Dept. of Architecture. Program in Media Arts and Sciences. Cambridge, MA, Massachusetts Institute of Technology.

Stock, O., M. Zancanaro, et al. (2007). "Adaptive, intelligent presentation of information for the museum visitor in peach." User Modeling and User-Adapted Interaction 17(3): 257-304.

Wakkary, R., M. Hatala, et al. (2005). An ambient intelligence platform for physical play. ACM Multimedia 2005, Singapore, ACM Press.

Wakkary, R. and M. Hatala (2007). "Situated play in a tangible interface and adaptive audio museum guide." Journal of Personal and Ubiquitous Computing 11(3): 171-191.

Cite as: