DOCUMENT RESUME

ED 482 046 IR 058 801

AUTHOR	Bearman, David, Ed.; Trant, Jennifer, Ed.
TITLE	Museums and the Web 2003: Selected Papers from an International Conference (7th, Charlotte, North Carolina, March 19-22, 2003).
INSTITUTION	Archives and Museum Informatics, Pittsburgh, PA.
ISBN	ISBN-1-885626-27-1
PUB DATE	2003-00-00
NOTE	260p.; Accompanying CD-ROM not available from ERIC. For selected individual papers, see IR 058 802-823.
AVAILABLE FROM	Archives & Museum Informatics, 158 Lee Ave., Toronto, ON, M4E 2P3 Canada. Tel: 1-416-691-2516; Fax: 1-416-352-6025; e-mail: info@archimus.com; Web site: http://www.archimuse.com/. For full text: http://www.archimuse.com/mw2003/ .
PUB TYPE	Collected Works - Proceedings (021) -- Non-Print Media (100)
EDRS PRICE	EDRS Price MF01/PC11 Plus Postage.
DESCRIPTORS	*Archives; Art; Information Technology; Institutional Libraries; Internet; *Museums; Program Development; Special Libraries; World Wide Web
IDENTIFIERS	Museum Collections; *Museum Libraries; Virtual Museums

ABSTRACT

This is the proceedings of the seventh annual Museums and the Web conference which took place March 19-22, 2003. MW2003 was the premier international venue to review the state of the Web in arts, culture, and heritage. The formal program consisted of two plenary sessions, eighteen parallel sessions, 35 museum project demonstrations, dozens of commercial exhibits, full-day and half-day pre-conference workshops, and one-hour mini-workshops combined with a day-long usability lab, a day-long design "crit room", and the Best of the Web awards. Participants were webmasters, educators, curators, librarians, designers, managers, and directors who work in museums, libraries, archives, historic sites, and the companies that support them. A contents table lists, in alphabetical order by contributor, the speaker, speaker's country, title of the paper, type (paper, workshop, demonstration, etc.), and whether the paper is online. Papers describe innovations and developments in programs, procedures and technology in museums and arts institutions including: "Integrating Databases with Maps: the Delivery of Cultural Data through TimeMap" (Ian Johnson); "Software Tools for Indigenous Knowledge Management" (Jane Hunter, Bevan Koopman, Jane Sledge); "Dublin Core: The Base For An Indigenous Culture Environment?" (Liddy Nevile, Sophie Lissonnet); "A Prototype Digital Library For 3D Collections: Tools To Capture, Model, Analyze, and Query Complex 3D Data" (Jeremy Rowe, Anshuman Razdan); "The More You Look the More You Get: Intention-based Interface using Gaze-tracking" (Slavko Milekic); "Re-assessing Practice: visual art, visually impaired people and the Web" (Caro Howell, Dan Porter); "From GUI to Gallery: A Study of Online Virtual Environments" (Stephen Lawrence Guynup); "Interfacing the Digital" (Steve Dietz); "Practicing What We Teach: How Learning Theory Can Guide Development of Online Educational Activities" (David T. Schaller, Steven Allison-Bunnell); "Evaluating the Authenticity of Egyptian Cartonnage Fragments: Educational Outreach in Search of the Truth" (Paul Marty, Kim Sheahan, Ann Lacy); "Focus your young visitors: Kids Innovation Fundamental changes in

digital edutainment" (Sebastian Sauer, Stefan Gobel); "Investigating Heuristic Evaluation: A Case Study" (Kate Haley Goldman, Laura Bendoly); "New Vision, New Realities: Methodology and Mission in Developing Interactive Videoconferencing Programming" (Patricia Barbanell, John Falco, Diana Newman); "A Rolling Evaluation Gathers No Moss" (Lee Anne Burrough, Lorrie Beaumont, David Schaller, Ethalinda Cannon); "The SEE Experience: Edutainment in 3D Virtual Worlds" (Nicoletta Di Blas, Paolo Paolini, Susan Hazan); "Make Your Museum Talk: Natural Language Interfaces for Cultural Institutions" (Stéfania Boiano, Giuliano Gaia, Morgana Caldarini); "Interactive Character as a Virtual Tour Guide to an Online Museum Exhibition" (Pilar de Almeida, Shigeki Yokoi); "Experiencing Art on the Web with Virtual Companions" (Ido A. Iurgel); "Using Cinematic Techniques in a Multimedia Museum Guide" (M. Zancanaro, O. Stock, I. Alfaro); "The State of the Art in Museum Handhelds in 2003" (Nancy Proctor, Chris Tellis); and "Designing Multi-Channel Web Frameworks for Cultural Tourism Applications: the MUSE Case Study" (Franca Garzotto, Tullio Salmon Cinotti, Massimiliano Pigozzi). An accompanying CD-ROM includes: a list of all the speakers at the conference and links to their abstracts, biographies, and papers (where available); an overview of the Museums and the Web 2003 conference program and links to abstracts and paper biographies; and the results of the Best of the Web 2003 conference (requires Internet connection). Author biographies are included. Most papers contain references. (AEF)

www.archimuse.com/
mw2003/

Museums and the Web 2003

Selected Papers from an International Conference

edited by
David Bearman and Jennifer Trant

IR058801

Archives & Museum Informatics
158 Lee Avenue
Toronto, Ontario
M4E 2P3 Canada

info@archimuse.com
www.archimuse.com

Museums and the Web 2003

www.archimuse.com/
mw2003/

Archives & Museum Informatics

Consulting, Publishing and Training for Cultural Heritage Professionals

***Informatics**: The interdisciplinary study of information content, representation, technology and applications and the methods and strategies by which information is used in organizations, networks, cultures and societies.*

Archives & Museum Informatics organizes conferences, workshops and seminars, publishes articles and monographs, and consults for cultural heritage organizations worldwide. For over 12 years, our educational goal has been to provide specialists in archives, museum and cultural and heritage informatics with timely and challenging opportunities for professional exchange and training. Our consulting services emphasize inter-institutional collaboration, strategic planning and standards-based solutions.

To be kept informed of coming events, conferences and new publications, contact us at:

Archives & Museum Informatics
158 Lee Avenue
Toronto, Ontario
M4E 2P3 Canada
phone: +1 416 691 2516
fax: +1 416 352 6025
email: info@archimuse.com
http://www.archimuse.com

Ask to be put on our mailing list.

www.archimuse.com/
mw2003/

Museums and the Web 2003

Selected Papers from an International Conference

edited by
David Bearman and
Jennifer Trant

Museums and the Web 2003
Selected papers from an international conference

edited by David Bearman and Jennifer Trant
Toronto, Ontario, Canada: Archives & Museum Informatics
ISBN: 1-885626-27-1

Printed in the United States of America

Contents

Introduction

Representing Cultural Knowledge

Interfaces

Design for Learning

Evaluation

Virtual Visiting

Delivering Digital Heritage

About the Authors

About the CD-ROM

Acknowledgements

Our sincere thanks once again to the many authors of papers for Museums and the Web 2003, who have demonstrated – by their writing – the value of a published literature documenting Museums and the Web. Their efforts help move all our activities forward.

Our thanks too to the many authors and presenters whose work appears on the accompanying CD-ROM and on the conference Web site. The Museums and the Web site can be found at http://www.archimuse.com/mw.html. Please visit it to find many more papers, links to the sites described, and full descriptions of the demonstrations and workshops presented at MW2003 and all previous Museums and the Web conferences.

This volume was copy edited by Sheila Trant.

Introduction

Familiarity Breeds Content: knowledge and effect in museums on the Web

David Bearman and Jennifer Trant, Archives & Museum Informatics, Canada

http://www.archimuse.com

Abstract

The papers presented at the seventh Museums and the Web Conference reveal that the technologies of the World Wide Web have – over the past ten years – significantly influenced the business of museums and changed how it is conducted. This introduction to the Selected Papers volume in print focuses on developments in knowledge representation, interface design, program presentation and learning theory, evaluation, the techniques supporting virtual visiting, and multi-channel information delivery. In the conference itself, an even wider set of issues were addressed. Throughout, what is clear is the speed with which a new communication genre is evolving and the radical impact that ubiquitous electronic delivery is having on a cultural institution type that, during its first two centuries of existence, was largely unchanged by technology.

Keywords: knowledge representation, interface design, online presentation, learning theory, virtual visiting, interactive, multimedia, Web evaluation, multi-channel information delivery, museums, museum missions, museum business functions, process re-engineering, electronic genres, new media

Introduction

Every year, as the co-chairs of the International Museums and the Web Conference, it is our pleasure to review hundreds of contributions from around the world to the burgeoning literature on museums and new media programs. And every year, when we select the papers that we feel will have the longest term impact on the field, we are delighted and surprised to tease out the prevailing themes of discourse and reflect on their significance to the field.

The Museums and the Web Conference (MW) itself is not just a context for academic discourse, though these papers can hold their own in any intellectual discussion: MW is a many ring extravaganza, with designers holding "crit" rooms, evaluators conducting usability labs, techies teaching mini-workshops on methods, professional debates in open forums, demonstrations by dozens of museums of their redeveloped sites and exhibits from dozens of commercial companies featuring their approaches to the field. It is above all an arena of practice, where what counts is what you can do, not abstract ideas that haven't been realized. Hence, these papers only begin to suggest the full scope of the meeting, a record of which can be found online along with all prior meetings since 1997. On the Web you can find the full text of all presentations, abstracts of Demonstrations and speakers' biographies (this and prior years linked from http://www.archimuse.com/mw.html).

Data, Content and Knowledge Representation

Before the Web, much museum documentation was "just data", raw information that told us what we had collected and where we put it. But in a short decade everything has changed. Now that museums and other cultural institutions are judged by the quality of the *experience* of their Web sites, the question of what you can do to present information in exciting and useful ways is paramount. Heightened user expectations have a direct bearing on the initial capture and subsequent management of the digital documentation of museums. If the name of the ancient city where parts of a burial urn was unearthed is recorded by the museum, whether it is stored in such a way as to make it possible to display it correctly and automatically in a cultural historical atlas, or search it with a GIS (geographic information system) interface, is central to its value as a museum asset. Ian Johnson's discussion of TimeMap and the Electronic Cultural Atlas Initiative show us the potential for this kind of view of museum collections. If the information respecting this burial urn is attributed to the various sources that could tell us about it and its meaning, including the ancestors of the people whose practices it represents, we not only have layers of professional and cultural meanings to impart. We can also respect sacred cultural meanings, reserving access to those entitled to share in them, as the metadata and software designs advanced by Hunter et. al show. If, in addition, we record data in a manner consistent with the Dublin Core metadata for discovery, sharing it with other museums may assist in the virtual repatriation of knowledge, as explored by Nevile and Lissonnet. If complex representations of the 3-D characteristics of the artifact are recorded, we can link it to related works (or other parts of the same work) in other collections and build a visual representation of the object that will enable future virtual visitors to have an experience which is as complete as (or more complete than) those coming in person. Such models, as described by Rowe and Razdan, are themselves complex knowledge representations, and it matters a great deal what rules are used in their preparation.

The four articles that comprise this section on knowledge representation bring home a number of crucial issues that museums are facing in the era of the Web – from the very first encounter with an object through all subsequent study and use, the museum must make a choice between capturing information in ways that will support future re-uses, many as yet unknown or unimagined, and methods that will end with the task at hand. Because of the significant cost of handling original objects and devoting scholarly attention to them, the "right' answer is totally clear: reusable content captured with explicit attribution, at a high level of granularity and in standard formats is ideal. But the re-engineering of the business processes in museums and the implementation of technical methods required to capture re-usable digital content are still major dilemmas. Furthermore, the very articulation of the problem – creating data today in ways that will be most usable and exciting tomorrow – exposes the heart of the conundrum: what will be usable and exciting tomorrow?

Interfaceless Interfaces

We struggle to present rich cultural experiences, to make the virtual environment a pure and simple extension of the real world without special methods or boundaries. Already, as we will see in articles on virtual visiting and multi-channel delivery in this volume, "on-line" is rapidly becoming a retronym, along with "wireless". Interface devices will, according to the authors of the papers in this section, become as quaint as the dial on a phone (or a phone itself) in a generation's time. Slavko Milekic is imagining and designing a future without interface devices. He brings the careful methods of the neurologist and the imagination of the artist to bear on systems which respond to us precisely the way humans respond – by gauging the extent of our interest and attention and providing us with feedback based on our gaze. Caro Howell and Don Porter push the interface to account for the sensory abilities of the viewer and bring visual art to life for those without vision. By designing an interface that does the impossible – in the case of the Tate, being a visual art museum to those who cannot see – they show us that interfaces can bring museums to an audience that couldn't use them before by enabling the interpenetration of the virtual and the physical worlds. They point us towards a future in which interfaces not only permit interaction with four of the senses – touch, smell, vision and hearing – but surpass the native limitations of those senses. Steve Guynup begins with the limitations of the current screen and mouse based virtual environments (SMVE's) and work he has done in 3D game design to suggest how the next-generation interfaces in galleries and the Web might need to feel, and moves from there to the inspirations of new virtual art for design paradigms. All art will never be digital – after all we are not leaving the physical world behind – but digital art is becoming as broad and legitimate an expressive medium as any other technique of artistic expression. As Steve Dietz , who has curated digital art exhibitions for the past decade, suggests the digital is now an artistic domain that bridges the world of visual arts, performing arts, language arts by forging new syntheses, crafting new genres, and enabling new kinds of experiences. As he explores the concrete challenges of presenting such art within the Web and in the gallery, we are introduced to concepts of interface that move us dramatically beyond our current frame of reference and challenge us to respond creatively.

Learning, not teaching

Together these visions of the way that we will relate to content – and it to us – begin to suggest ways in which the museum, and its collected knowledge, can be drawn into a dialogue with inquisitive minds everywhere. Gone (after but a few years) is the fixation with connecting (wires) to the Internet, and of special applications (instructional software) to teach set curricula in fixed loci. Instead, we are evolving anytime, any-question storytelling, exploration, way finding and meaning-making. David Schaller and his colleagues explore the constructivist learning theories that are the scaffolding on which numerous successful learning environments have been built, and suggest how learning theory can provide criteria by which to assess the value of future efforts. Their grounding of concrete museum application designs in learning theory reveals both how much we have to discover about learning, and why it is useful to use what we do know in constructing Web programs. Paul Marty and his colleagues take a specific problem in knowledge and skill building – the recognition of authenticity – and demonstrate how Web-based learning situations designed by the museum can bring learners into direct engagement with objects. This assignment both tests exploratory methods that develop the skills to assess what is real, and teaches students how to exploit what is virtual and metadata only. Where all evidence is presented as virtual representations, these skills will increasingly form the foundation both of computer-literacy and museum-literacy. At the root of the success of these didactic experiences is that they are, nonetheless, very engaging. Sebastian Sauer and Stefan Göbel explicitly address the question of what makes "edutainment" and how to mix pedagogic aims with the fun of play, the excitement of the unknown, and the comfort of a knowledgeable companion. Their Monuma pocket device for exploring museums, and the Telebuddy avatar/familiar (also described at MW2003 by Anja Hoffmann and Stefan Göbel) are (still primitive) prototypes for the future of edutainment. They provide us methods to test the ways that children, and adults, learn from accompaniment by a knowledgeable (quite cuddly) presence.

Evaluation

Increasingly, we are learning how to determine if what we are doing on the Web works and finding that, often, it does and that it certainly can. The continually growing literature on evaluation and its increasing sophistication encourage and validate investment in assessing our activities. In the first of what we hope will be a series of self-conscious evaluation reports, Haley-Goldman and Laura Bendoly define a methodology to rigorously compare the heuristic evaluation methods – often favored by museums because of their cheapness and speed – with other methods of evaluation. While the study is still in progress, the framing of the question with respect to the record of heuristic evaluation in other domains is already an important contribution, and the methodologies suggested for the longer-term cross-method testing will contribute to evaluation of museum Web activities in the future. Barbanell and her colleagues have conducted a major collaboration with teachers and museums engaged in video-conferencing in conjunction with the Web across the State of New York. A substantial part of their effort has been directed to building a framework for assessing such technological, pedagogical and institutional innovations. Concrete evaluation results from these projects reported elsewhere at MW2003 validated the video-conferencing initiative. Burrough et. al. have implemented an interactive design experience in which kids and experts interact in real-time, and in virtual space, to meet a concrete challenge in engineering. Kids Design Network involves teachers as co-learners and adults as mentors. This challenging multi-model program has become the object of continual evaluation and feedback refinement. Different evaluation methods were used during the course of the implementation, and the development process was engineered with distinct points at which to take findings into account. By subjecting the entire venture to continuous, multi-method evaluation and feedback/redesign, this group demonstrates how the genre of interactive learning can best evolve, and reports on an important instance of using the Web, remote expertise, and classroom problem-setting to construct a compelling educational experience.

Virtual Visiting

Many of the themes of the *Selected Papers* volume and the conference converge in this chapter. Robust knowledge representations and novel interfaces meet the learning environment in a successfully synergetic application metaphor – the virtual guide. Not surprisingly, these applications also pose new challenges to evaluation methods and (presaging later papers) demands for multi-channel delivery. The numerous papers given at this year's conference on the virtual exhibits, the virtual visitor experience and storytelling and narration come together in the effort to craft lively virtual companions that can engage in entertaining, knowledge-based discourse with visitors to virtual museums. The technical, social, engineering and ontological issues posed by this application have been finding their way into more and more Museums and the Web papers each year; the progress that has been made in a few years is astounding, though the distance that we still have to travel (carefully delineated now) remains vast.

The first two papers in this section extend the pioneering work (now four years old) at the Science Museum of Milan. Di Blas and her colleagues have focused on the problems early implementation exposed in the virtual environment and visitor orientation. They are building a new generation of guide system that incorporates yet a more complex discourse structure, and connects to the technology of "chat" so popular in on-line spaces, especially with the young. Gaia and his colleagues push the boundaries of 'chatbots' towards virtual guide technology, offering a blueprint for the levels of problems to be resolved: obtaining appropriate knowledge; enabling written input and response; enabling spoken response, followed by spoken input; and ultimately connecting the bot's knowledge-base to the museum knowledge-base and beyond. These progressively more challenging stages in deploying such a bot document the research program museums face before successfully deploying this technology. In a simulated guided tour, Almeida and Yokoi have made considerable progress in narrowing the domain of questions, documenting the structure of the tour guide storytelling, and analyzing the means for linking databases in order to create a knowledge-base about concrete objects and contexts, and about the structure of discourse genres. Knowledge-base structuring and interaction, from written queries to spoken responses, frame this exploration. Finally, Iurgel has focused on the affective interaction important in art appreciation, and designed characters that are optimized to respond in ways that people might perceive as friendly and empathetic. The breadth of this research, no less than the fact that the researchers span the globe from Japan to Germany and Israel to Italy, reveals the vast interdisciplinary challenges involved and the ways in which this application arena sits at the heart of the new media revolution.

Multi-Channel Communications

What, we might reasonably ask, will be the future of museum Web activities, or of the Web itself? It is likely not to be very recognizable to those seeking computers or even monitors and expecting data to flow over wires to pre-determined places. The outlines of the webbed information services of the future are only now coming into view, but already it is clear that we will carry with us our own personal devices, that the information will follow us (or come to where we are), and that the form of the communication will not be broadcast, nor narrow-cast, nor point-to-point in the way that telephonic communication has been implemented. The model, rather, is of an information ecology, in which the environment is an organism that responds to our presence by delivering (to a broad range of receivers simultaneously), personalized responses (narrowly cast), that are delivered to any sensors we might be carrying or near (many points), all of which can interact to provide a variety of fully surrounding sensory/information experiences. The dashboard of our vehicle, our cell phones, the sidewalks we walk on, the walls we pass, and perhaps the clothing we wear, will respond with image, sound, and possibly tactile sensations, to inform, entertain, and engage us. And we will reply, in speech, with gesture and with touch, to stimulate, direct and structure the experience.

And how close have we come to the Web of tomorrow? Zancanaro et al. report on how to explore a single work of art hundreds of times the size of their PDA in a way that can take us into the details of the scene. With in the confines of the PDA screen, they suggest how a knowledge of cinematographic techniques and story-telling structures permits the construction of an engaging dialogue. Zimmermann et. al focus on the information brokering tools needed to mine the knowledge-bases of museums and locate the facts that form the core of personalized chat. They then link these methods, using talk as the interface, in an environment alive to actions and locations. Proctor and Tellis report on the state of handheld tour devices and detail the next generation of this rapidly developing technology. Their report on user surveys points out directions for future development. Finally, Garzotto et. al explore the museum implications of the explosion of delivery channels. An engineering framework for multi-channel delivery and a set of scenarios for user experience in different receiving environments are defined, and the requirements they pose for museums outlined. Finally, the practical task of constructing different interfaces to the same data, personalizing interactions, and creating guided (situated knowledge-based) strategies are explored.

Concluding thoughts, beginning ideas

Several major themes that have been present throughout the Museums and the Web conference series continued in evidence at the 2003 conference, but are not fully reflected in this book. (All the conference papers are available on-line at http://www.archimuse.com/mw2003/ and can also be found on the CD-ROM accompanying this volume.) The challenge of understanding our audience, and particularly the needs of children, continues to engage educators and public outreach professionals. Integrating Web resources within the physical exhibition space occupies exhibit designers. The problems of digital content management over the life of institutions are a focus of studies and tools presented by information technologists. As in the past, it was clear that the Web is forcing major changes in business practices on the museums and that it is engaging nearly all the knowledge workers in these institutions in rethinking their methods and the very purposes of their institutions.

In a few short years we've come a great distance exploring the role and reality of museums and the Web. We've moved from simple brochures to collections catalogs, beyond e-commerce and to interactive, on-line programs with targeted audiences and tie-ins to the physical museum. But the pace of change is increasing, and the near future holds a Web unlike the space we've been confined to for the past seven years. No longer tethered by content developed for another purpose and time and captured in legacy computing systems, we are challenged to create digital content as we work in ways that are re-usable, interoperable and sharable. Then we can seek to implement systems that make content nearly ubiquitously available with no apparent interfaces, that communicate with visitors and passers-by in natural language, and that respond to their affect and attention. We are challenged to script our discourse to convey knowledge in the course of entertaining and engaging conversations and lively social experiences.

In a decade, will the resulting institution look like a "museum"? It will, if a museum is a place where objects that represent culture and humanity, natural and physical diversity, are cared for and interpreted. It will if our experiences of museums on the Web confront us with novelties and leave us with understanding. But it will be a different place than it was in 1990, or even in 1997, when we first erected virtual billboards to announce the opening hours and location of the building.

David Bearman and Jennifer Trant
Co-Chairs, Museums and the Web
March 5, 2003
Toronto, Ontario, Canada / www.archimuse.com

Representing Cultural Knowledge

Integrating Databases with Maps: the Delivery of Cultural Data through TimeMap

Ian Johnson, University of Sydney Australia

Abstract

TimeMap[1] is a unique integration of database management, metadata and interactive maps, designed to contextualise and deliver cultural data through maps. TimeMap extends conventional maps with the time dimension, creating and animating maps on-the-fly; delivers them as a kiosk application or embedded in Wweb pages; links flexibly to detailed content in Wweb pages and databases; connects to a wide variety of data sources, including textual databases and scanned historical maps, situated anywhere on the Internet; and allows locational data to be captured in a Wweb browser. This paper presents an overview of TimeMap, focusing on the flexible data model and metadata-based methodology which allow TimeMap to integrate seamlessly with existing databases and fulfil a wide range of application needs from museum kiosk to digital cultural atlas. Example applications include the Sydney TimeMap museum kiosk, MacquarieNet on-line encyclopaedia, Fairfield multicultural tourism Web site, UNESCO World Heritage Sites database, and a map animation of the Khmer Empire. The paper also addresses some of the issues in handling complex datasets effectively on low bandwidth connections, and shows how TimeMap has addressed these through on-demand data loading, server-side filtering, dynamic map legends and dynamically constructed hotlinks.

Keywords: Time, GIS, Web mapping, Database integration, Java applets, e-Publishing, Animation, Electronic Cultural Atlas Initiative, TimeMap

Introduction

This paper is about the integration of maps and databases, and the use of maps as a front end to rich cultural information. While static maps – and, to a lesser degree, purpose-built animated maps and interactive multimedia with fixed content – are commonplace in museums and on the Web, the use of interactive maps built on-demand from content-rich databases and the incorporation of time as an essential element of historical maps are altogether less common. This paper will briefly review the strengths of interactive time-enabled maps and obstacles to their wider use, and outline the methods developed by the University of Sydney TimeMap Project (www.timemap.net).

Why Maps, Why Databases?

Few people would deny the utility of maps as a means of contextualising cultural data, and fewer still would deny the utility of databases in keeping track of the vast amounts of information we all collect. Yet maps and databases are rarely integrated outside the Geographic Information System (GIS) community.

While the maps we see in museums, in print, on TV and on the Web are increasingly constructed using GIS – the integration of databases, maps and spatial analysis – they are most often presented as static images (animated weather charts being a notable exception). Web-based mapping is becoming common for geographic and environmental data, but has been relatively little exploited for cultural and historical data. On-line databases are commonplace, but most searches are still text-based, and results are generally displayed as lists even if leavened with images or multimedia.

Yet maps not only provide an effective means of delivering cultural data (Johnson 2002:1), but also provide a means of displaying lists of resources[2] in a more engaging, coherent and comprehensible fashion, and an unambiguous method of searching for relevant information. As physical beings we routinely use mental and physical maps –whether going on holiday to an unfamiliar place or simply fetching a spoon from the kitchen drawer – to understand our surroundings and to find things of relevance to our needs or interests. In *On the Internet*, Dreyfus (2001:21) observes, "Our embodied concerns so pervade our world that we don't notice the way our body enables us to make sense of it". Egenhofer points out the weakness of textual searches as providing "no support for any deeper structures that ... people typically use to reason" (2002:1). Time and place are fundamental to our understanding of the world, the items it contains, the way things came

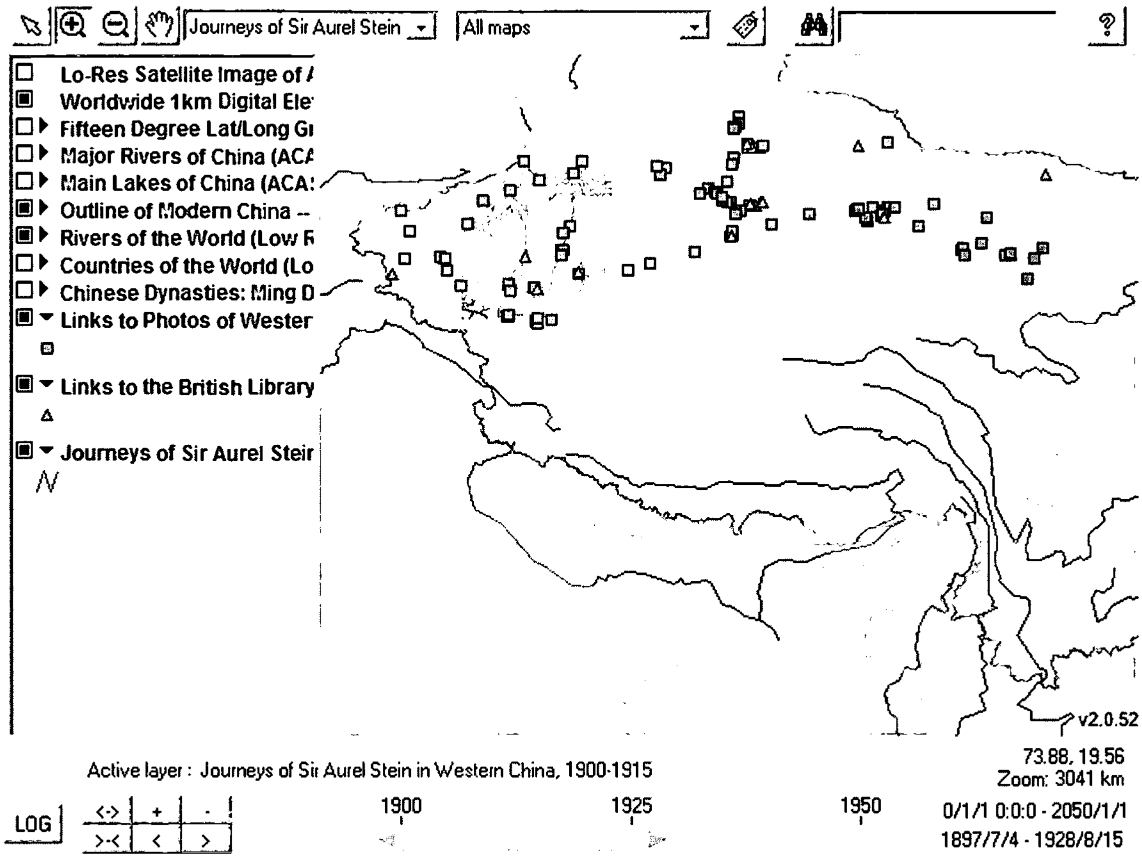

Figure 1:Time Selection Bar

about, our immediate concerns, and our future trajectories.

Maps provide an interface which is more consistently understood, less ambiguous, more portable across cultures and more grounded in our experience than classification, search terms and ordered lists. The latter aim to simplify the complexity of the real world, but they are only useful to a constrained social and cultural group. On the other hand, we need only to be able to implement definitions of latitude and longitude – scarcely an issue today – to bring all spatially-located information into a common framework and to agree on a spatial search. While methods of representation, such as projection, symbology and choice of data, may substantially condition the message transmitted (see Monmonier 1996: 94-107 on the use of maps for propaganda), maps also provide a rich and concise means of delivering large volumes of information from a database back-end.

By building explicit programmatic methods for map construction on top of a database, rather than handcrafting a multimedia resource from a fixed set of content, one has the flexibility to expand, enhance, redesign or even repurpose the package without starting again from scratch. Our aim with TimeMap has been to develop generalisable methods rather than specific solutions, so that a wide variety of interactive map-based kiosks and Web sites delivering historical data in spatial and temporal context can be developed with little or no programming.

What is a Time-Enabled Map?

There are a number of conventional mapping methods for representing change through time which can be seen in historical atlases such as the *Concise Atlas of World* History (O'Brien 2002) or the *Atlas of World Archaeology* (Bahn 2000): period maps, time symbolisation, change maps and annotation. These methods are static and have problems displaying very complex or detailed overlapping information, particularly where features undergo many small changes or change continuously through time (administrative boundaries and the extent of empires are examples of such temporally and spatially complex objects).

Date Stamping and Filtering

Computers allow us to generate maps on demand so that we can set a time range and more-or-less instantly see what falls within that range. TimeMap implements this method of time-enabled map display through a time selection bar (Figure 1) and filtering of the source data to display only those features which fall within the time range selected[3]. TimeMap will also step through a sequence of time ranges to display through-time changes on the map interface or to generate an animation file which can be played separately or embedded in a Web page.

Many datasets can be handled with simple date stamping along these lines, and such date stamping is often already available within databases of historical information as start/end dates, a simple median date or an event date. Simple date stamp filtering works well for items which occur at a specific time identified by a date or date range (such as historical photographs), or exist unmodified over a period of time (such as particular administrative boundary configurations), but is less effective for items which change continuously during their period of existence (such as empires or cultural areas).

Snapshots and Interpolation

In practice, our knowledge of historical data tends to be episodic and partial, so that we know the position, extent or attributes of a map feature – a "snapshot" of the feature – at various (arbitrary) times, but do not have specific information on what happens in between[4]. We can make some estimate of intermediate conditions through interpolation of extents or attributes between the available snapshots (Figure 2).

In the absence of other information, such interpolation would normally be linear between the known snapshots. However we may have enough information about the nature of the phenomenon to record a qualitative modifier, such as "increasing pace", or even a quantitative modifier which could be used to improve the interpolation. Furthermore, the dates of the snapshots may be subject to uncertainty, as may the accuracy of the extents or attributes recorded, and all of this information should really be recorded and used in the generation of interpolations. I refer to this method of recording spatio-temporal information as a "snapshot-transition model", and refer to the individual observations of features at arbitrary times as "instances" (Johnson 1999).

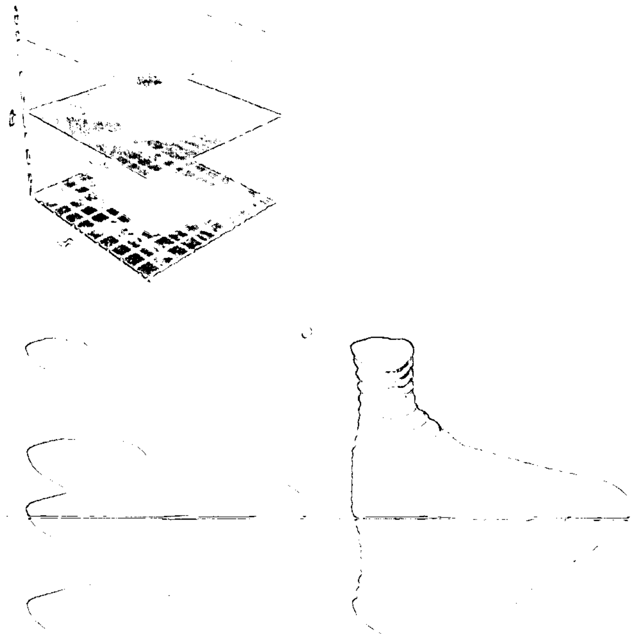

Figure 2: The snapshot-transition model of a feature (viewed in a 'space-time cube'): snapshots (left) and interpolated sequence (right)

Model-driven Interpolation

While modifications to the interpolation of features recorded using a snapshot-transition model may produce a better match with what actually occurred, such interpolations are still just mathematical transformations which take no account of what actually happened, and the degree to which they match reality can never be known. The only way to generate interpolations which approximate reality, within some quantifiable error, is to model the processes which resulted in change, taking account of factors such as landscape characteristics, social organisation, economics, communication and specific events. The model might then be tested against actual observations - the snapshot observations.

Such an approach requires far more knowledge and far more data than is generally available for historical periods, but should not be excluded as a means of refining interpolations. The development and refinement of such a model may also be a valuable method for better understanding and interpreting the phenomena mapped, as well as delivering visualisations for explanation and education.

Map Animation

Map animations generated either by simple geometric interpolation of instances (tweening) or through a model-driven interpolation provide a powerful and engaging means of communicating information about mappable historical phenomena – distributions of sites and cultural domains, trade, exploration, spread of empires, military campaigns, urban growth etc. Unlike simple time date stamp filtering, animation works best where features DO change continuously (although they can also work effectively where there is patterning in the appearance and disappearance of substantial numbers of point features).

Map animation requires the interpolation of known spatial extents to generate intermediate 'frames'. For bitmap animation formats, such as AVI or MPEG, all intermediate frames must be generated, and the resulting files are very large even with compression. Our early small animations ran to tens of megabytes.

Vector animation formats such as SVG and SWF (Flash) can be rendered by a player which interpolates intermediate frames through tweening. This closely parallels the snapshot-transition data model adopted by TimeMap, so we can export an SWF file for tweening and animation by the Flash player, which can also be imported into Flash MX for refinement (see Jesse, 2001 for a similar approach). Because frames are built on demand and based on vectors, the size of these Flash animations depends on the amount of temporal and spatial detail, not on the playing time of the animation, and good quality zoomable animations running for several minutes are typically less than 1 MByte. Example animations are available on the TimeMap Web site.

Barriers to adoption

There are cogent reasons why the integration of maps and databases has not occurred widely. The high cost of software and data preparation and the need for technical expertise and facilities have been major hurdles. These technical and financial impediments to integrating maps with databases have by now been largely eroded, and most people have come in contact with database-generated interactive maps (without necessarily recognising the essential back-end behind the visual product).

The remaining barriers to widespread use of maps as a front-end to databases lie in awareness of non-demanding solutions and in seeing their relevance to providing more informative and more engaging interfaces for both search and display. One of our aims with the TimeMap project is to make database-driven map interfaces available, not as costly custom solutions, but as flexible search and enquiry systems which should be (almost) as easy and natural to use as images in a Web page.

Interactive Mapping Solutions

Web-based mapping is a growing field with many players and potential applications (for an overview see Kraak and Brown, 2001). At the lower end of the market there are any number of Java mapping applets–many available as freeware – which can be used to create interactive maps embedded in Web pages. Examples include ALOV[5], GAEA, GeoTools, JMapplet, JShape, lava, MapFactory and many others which can be located through a Web search. These applications – at least in their basic form – typically provide closed systems which download GIS data files stored with the applet, hold these in memory and allow various standard operations such as zooming, panning and enquiry. Like a Web page, you get all or nothing. These single-tier approaches are simple, but limit flexibility and entail heavy network traffic for rich datasets.

A number of these applet developers – along with digital library initiatives and mainstream commercial vendors – have also pursued a multi-tier approach in which the client can request data as required from a middle-tier server application. This approach gives greater flexibility to search for data through a clearinghouse, to add and remove layers, to pull data from multiple sources and to download data incrementally as a map is panned or zoomed. By isolating the complexities of data access and error handling on the middle tier server, the client applets can be simpler, smaller and more reliable, while network traffic is minimised.

The pre-eminent commercial multi-tier system today comes from ESRI Inc. Their Geography Network clearinghouse indexes (among other things) datasets served by their Internet map server (ArcIMS) to a variety of ESRI clients of various 'weights' (ArcReader ArcExplorer, ArcView, ArcGIS). Other vendors offer similar combinations of index, server and alternative clients.

Beyond these vendor-specific solutions, however, lies the potential for interoperability based on the Open GIS Consortium (OGC) standards (see www.opengis.org) which provide well-defined ways of finding, requesting and transferring data between data sources, catalogues, servers and clients. We should look beyond the traditional approach of building closed systems for specific applications – often developed independently, from scratch, as one-off applications – towards the building of systems which use adaptable tools and draw on external sources of rich data, such as mapping frameworks or institutional collections, and integrate them into new applications through standards-based interoperability.

TimeMap

TimeMap is intended to contextualise and deliver cultural data through maps. It is built around metadata stored in an SQL server DBMS, which allows data to be pulled in from servers anywhere on the Internet. To that extent it is similar in concept, although not in implementation, to commercial multi-tier systems. The biggest distinction from currently available commercial solutions lies in its explicit handling of the time dimension, informed by its roots in Humanities research (Johnson 1999).

The implementation differences lie particularly at the metadata level. Rather than a fixed set of metadata elements hard-coded into the software, TimeMap is a soft-coded system: the element definitions are recorded in the central server DBMS and can be consulted and used by each of the TimeMap programs. The definitions in the database can be easily updated; e.g. by changing their name or description or by adding additional schemes or permitted values, and these definition changes will

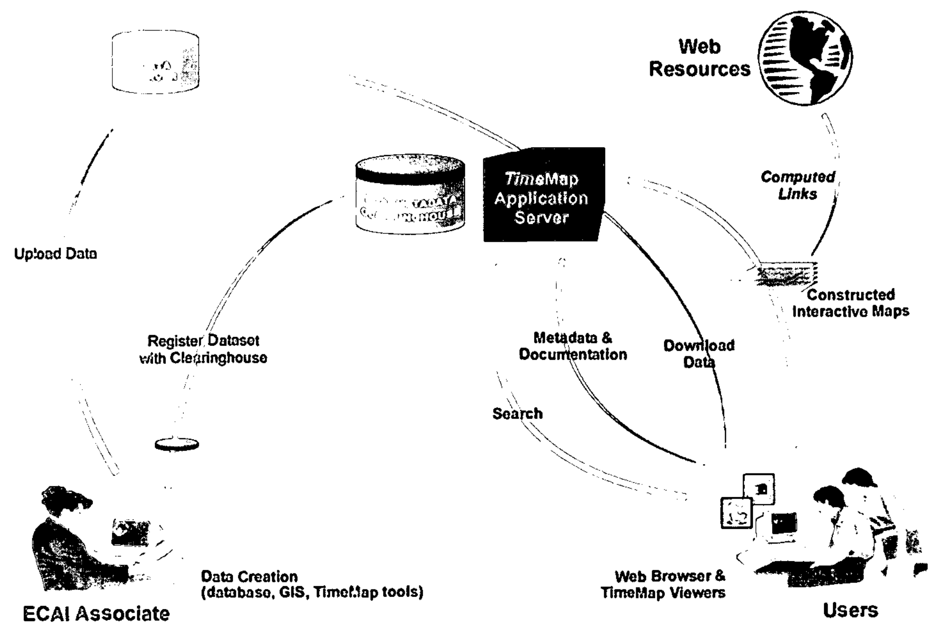

Figure 3: Structure of the TimeMap/ ECAI application

be reflected in the software without reprogramming. The metadata is not limited to simply describing the data, but is used to configure the software to the data, rather than requiring the data to conform to the software.

At the core of the TimeMap system is a central, Web-searchable index of datasets[6], the metadata clearinghouse developed for the Electronic Cultural Atlas Initiative (Lancaster & Bodenhammer, 2002). Datasets are spatially-referenced vector (point, line or area) data or raster images (scanned historical/ modern maps, aerial/satellite images, digital elevation images) which may be located anywhere on the Internet and be in a variety of common formats[7]. Figure 3 outlines the relationship between the different components of the system.

TimeMap allows datasets registered in the ECAI clearinghouse to be combined and overlaid in interactive maps ("MapSpaces"), which are set up using the TimeMap Windows program (TMWin). TMWin will publish MapSpaces to the clearinghouse and they can then be embedded in a Web page as a fully interactive Java applet (TMJava) with a few lines of HTML. TMWin can also be used stand-alone as a MapSpace viewer or kiosk application.

TimeMap in Practice

TimeMap is made up of a number of components which combine to deliver a complete electronic publication system for interactive Web-based or kiosk maps. The system may be best understood through a description of these components.

Serving data

Preparing datasets for use by TimeMap can be as simple as placing a JPG image or an ESRI Shapefile on a Web server and entering metadata to register the URL in the clearinghouse. TMWin, the Windows version of TimeMap, provides wizard-driven creation of metadata, simple spatial registration of scanned images, creation of Shapefiles from spreadsheets, and loading of GIS datasets into a public domain SQL server (e.g. MySQL, available on many ISPs and university servers). This brings the creation of fully capable, Web-accessible spatial databases within easy reach of non specialists at little or no cost.

The ECAI Metadata Clearinghouse

The ECAI metadata clearinghouse was developed as a central metadata reference for Internet-accessible datasets which might be accessed by TimeMap (or other) software. Currently there is only one generally accessible clearinghouse, the Electronic Cultural Atlas Initiative (www.ecai.org) clearinghouse, but the software can connect to other clearinghouses (and is used with local XML versions for standalone applications).

The clearinghouse records three types of metadata (for more detailed discussion of these metadata categories, see Johnson, in press):

- **descriptive metadata**, based on Dublin Core, used to search for relevant datasets
- **connection metadata** – host, server type, passwords etc – to tell software how to access the datasets
- **structural metadata** (semantics) to tell software how to form appropriate requests to the data server.

The clearinghouse can be updated through Web-based forms (implemented in Perl and Python) or through the metadata editing and registration function of the TimeMap Windows software, TMWin .

The Middle Tier

The TimeMap clients (TMWin and TMJava) connect to individual datasets through a middle tier Java servlet which isolates data access, translation and authorisation control in a single easily controlled location.[8] Multiple middle tier servlets can be implemented to scale the application, but so far the need for a mirrored or distributed index has not been felt.

e-Publication Toolkit

TMWin provides a complete solution for creating, serving and displaying interactive maps, including

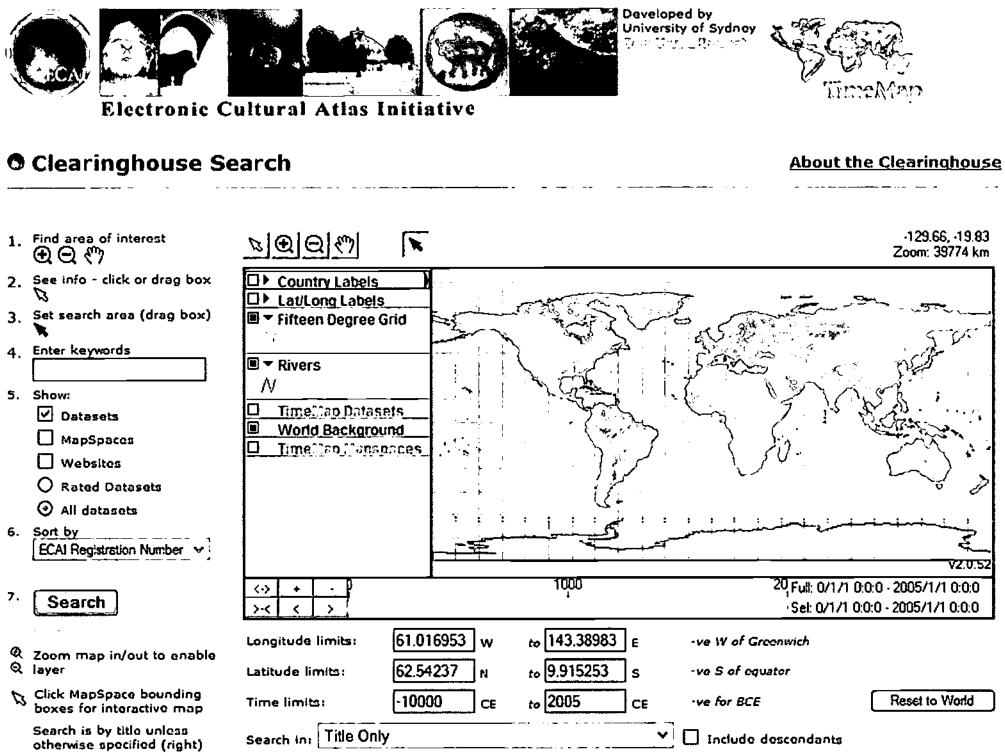

Fig 4: Clearinghouse search.

- metadata creation, editing, upload (registration) and download
- upload (data pump) of local spreadsheet or vector data to an SQL server database
- download of data from remote servers and local caching for offline use
- construction and upload of XML files defining an interactive map (termed a MapSpace)
- display of MapSpaces in fully-functional or restricted mode, for research, CD-ROM or kiosk applications.[9]

Client Applications

The Java version of TimeMap (TMJava) is used to render the MapSpaces created by TMWin and registered with the clearinghouse (or served as files from a Web server). The clearinghouse can be searched through a Web interface (Figure 4) which includes a TMJava map built from the clearinghouse database. Datasets selected through this search can then be combined and rendered as a TMJava map (Figure 5).

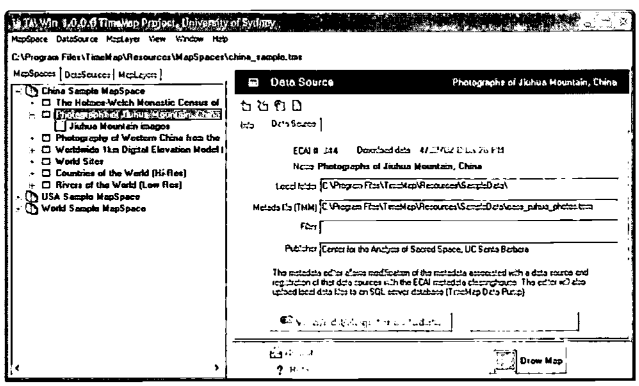

Figure 5: TimeMap Windows application (TMWin): MapSpace manager

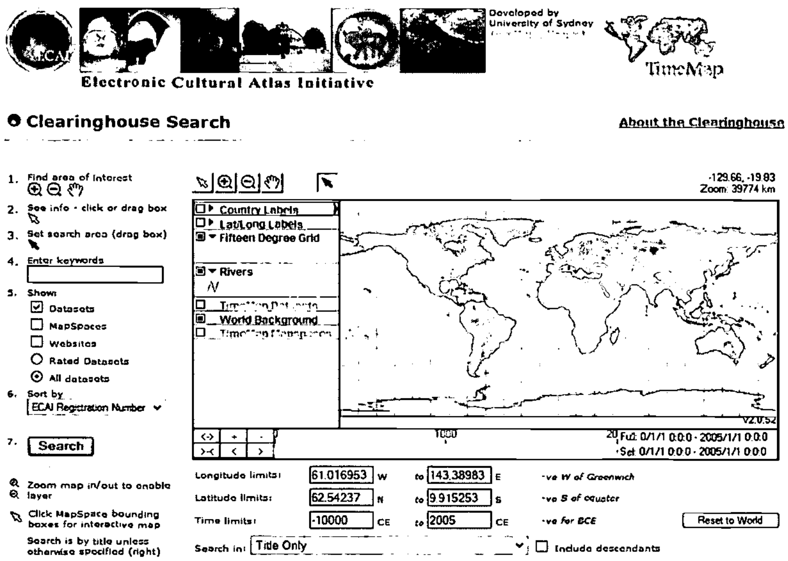

Figure 6: Electronic Cultural Atlas Initiative Clearinghouse search (the map is an interactive TMJava applet generated from the clearinghouse)

TMJava will embed a MapSpace as an interactive map in any Web page, with most of the symbology and behaviours associated with the Windows version. MapSpaces can also be displayed directly in a TMJava applet simply by supplying the registration ID to a special URL, as in the following example:

> http://www.timemap.net/clearinghouse/html/alov.cgi?id=329

If the default layout of the applet is not suitable, the applet can be extensively customised through XML layout files and applet parameters, allowing exposure of different levels of functionality and changes to the look and feel of the applet to fit with existing Web page designs. The applet can also be controlled by JavaScript, allowing complete custom applications to be built without modifying the core software. Several examples of default and customised TMJava maps, along with simple embedding instructions for user-designed Web pages, can be found on the TimeMap Web site.

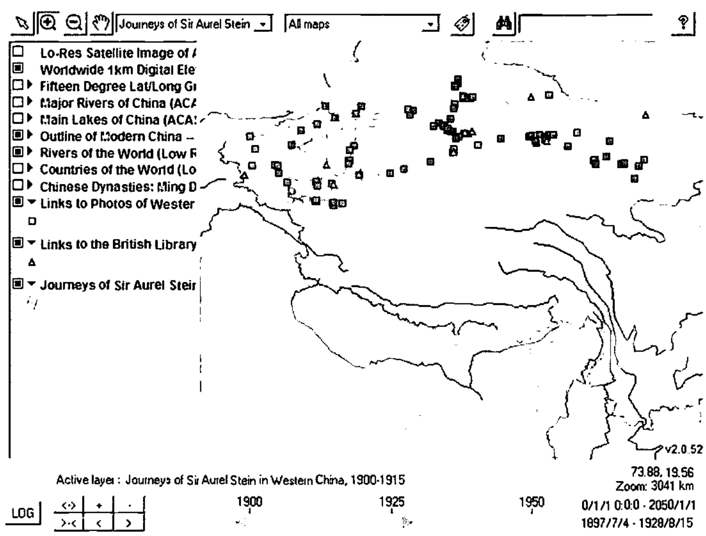

Figure 7: Search results from an ECAI clearinghouse search rendered in TMJava

Kiosk applications

It was always intended that TimeMap could function as a kiosk application, and our first kiosk (Figure 8) was installed at the Museum of Sydney in November 2000 (Johnson & Wilson 2002);. This kiosk is stand-alone, using resources generated from a central database and then copied to the kiosk computer itself.

Work since then has aimed at creating a distributed kiosk system which draws data dynamically from a central server. The current Museum of Sydney system was implemented in TMJava, which renders MapSpaces built on-the-fly from selections made in a ThinkMap™ application connected to the central

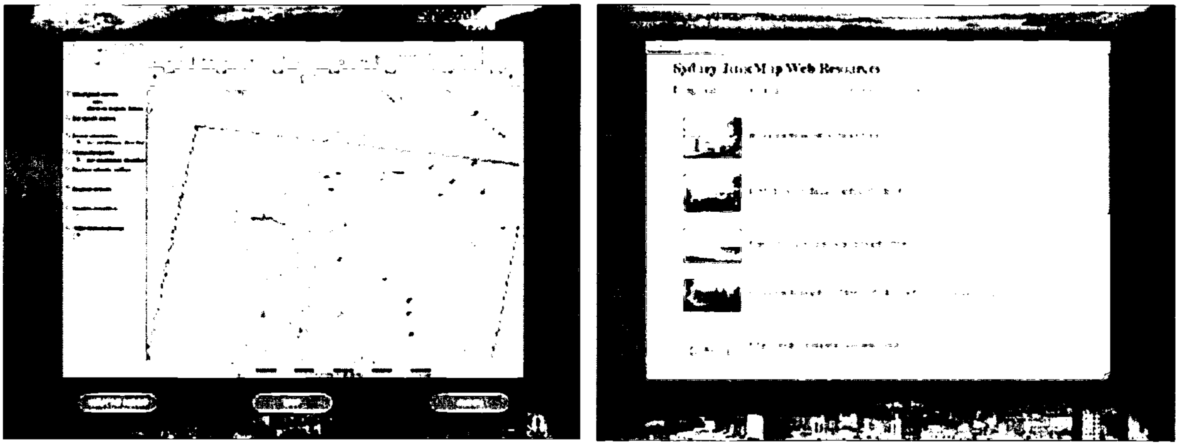

Figure8: Museum of Sydney dual monitor pilot TimeMap kiosk (Nov 2000)

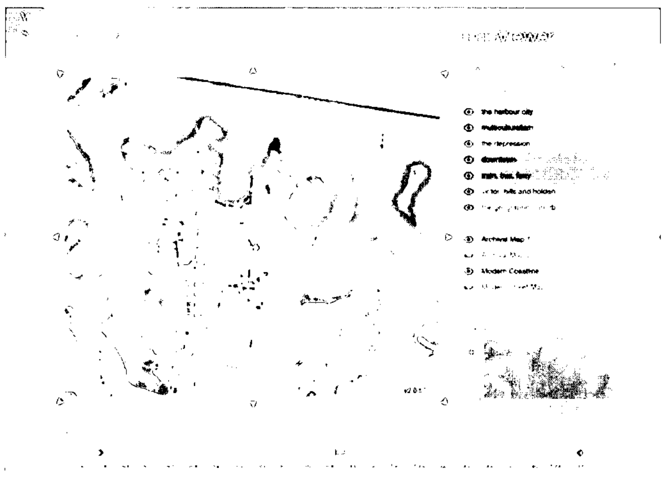

Figure 9: Museum of Sydney single monitor TMJava kiosk (Dec 2002)

server database. Specific customisation was carried out to match the Museum's screen designs and desired workflow (Figure 9).

Hot-linked content

One of the strengths of interactive maps is the ability to use the map as a portal to more detailed information This additional information may be the attributes (database fields) associated with the objects displayed on the map, or more detailed information such as database-generated Web pages or multimedia resources. Most interactive map applications now provide a means of hot-linking objects on the map to a disk file or a URL specified in a field in the attribute data, and in some cases support macros to generate a URL dynamically from information such as attributes and text.

TimeMap stores URL templates in the metadata attached to a dataset, which are then completed with attribute data when objects are selected on the map. This allows TMView and TMJava to easily hot-link to database-generated Web pages and systematically-named file resources, without the need to write macros.

TimeMap also handles hot-links from multiple layers, multiple hot-links per object and multiple object selection. When multiple objects are selected, a Web page of hyperlinks is generated from an HTML template.[10] The templates allow the list of hyperlinks to include thumbnail images (as illustrated in Figure 6, right-hand screen).

Bandwidth issues

Bandwidth is a critical issue in delivery of Web-based resources. Owing to the potential size of GIS datasets, many Web mapping applications process the data server-side and deliver maps as JPG images. Each operations on the map returns to the server for a new image. This "thin client" approach can provide low-bandwidth solutions which side-step problems with browser compatibility, but it severely limits interactivity. The TimeMap middleware can generate map images embedded in an HTML template if required for applications with minimal browser requirements or interaction.

At the other end of the scale, standalone map applets tend to download all the data to memory, even though much of it may never be displayed, and manipulate the data on the client – a "thick client" approach. While interactivity is enhanced, the bandwidth consumed may be out of proportion with the detail displayed.

TimeMap employs a mid-weight client which processes and filters data server-side, so that only the data needed to display the map is transferred across the network Several different approaches are adopted to minimise bandwidth. First, as layers are switched on or the map is zoomed or panned, the client works out what additional data it needs and sends incremental requests to the server. Secondly, map layers are scale sensitive, so that high resolution layers can be set to switch on only when the map is zoomed in, so that only a small part of the dataset is ever requested. Thirdly, MrSID image layers are filterable and multi-resolution, so the client requests images which cover just the visible window at an appropriate resolution. Finally, data is efficiently compressed before transmission.

Conclusion

In this paper I have attempted to summarise the TimeMap methodology, from its simple data model, core database and underlying metadata structure, to the practical implementation of e-Publication and Web map delivery tools. The implementation of TimeMap provides museums and other cultural organisations with a low cost, easily customised method of delivering rich content from databases, integrating external sources of data and taking advantage of spatial information and the power of maps to contextualise information.

This approach contrasts at one end with user-hostile text-based searching of catalogue databases, and at the other with user-friendly multimedia systems designed for a specific and limited purpose. Some may doubt that content-rich systems are appropriate to a museum environment, where quick turnover on scarce computer terminals may be desired. But as hardware becomes more ubiquitous, and the outreach of museums to the Web increases, the arguments for content-rich, Web-accessible multimedia generated from museum resources will increase.

TimeMap is a first, or perhaps second, generation system for delivering content-rich multimedia through Web browsers (whether on-site or viewed at home). In a few short years it will look primitive as 3D virtual environments become commonplace. But if it is built on sound database and metadata principles – as we hope it is – we can update the programming to integrate these new methods, rather than ditching the dated information architecture and software of a monolithic multimedia system. Our audiences may be becoming more sophisticated users of computers at break-neck speed, but that doesn't mean that their underlying questions and motivations are evolving like fruit flies: good structure and data today will be good structure and data tomorrow.

Notes

1 TimeMap is a registered trademark of the University of Sydney

2 Most collections of tangible or intangible objects have some inherent spatial distribution, whether it be provenance, reference or current location.

[3] TimeMap handles both single date stamps and range date stamps, with available precision down to seconds and back to geological time. Date stamps can be recorded either as numeric values in years or as ISO 8601 format dates (e.g. 2003-01-15 12:57:21). BCE dates are represented as negative values. The date stamp fields are indicated through metadata fields giving the type of date stamping and the field names for the date stamp fields – field type is handled automatically.

[4] Note: some authors refer to "snapshots" when all features are recorded at a given time – as in an aerial photograph or map – rather than snapshots of individual feature at different times according to availability of data – as in occasional observations of landscape conditions.

[5] The current version of Java TimeMap, TMJava, is an extended version of the ALOV map applet. An earlier version was based on GeoTools.

[6] The index is implemented as an SQL server database with supporting web management and access tools in Perl, Python and Java, and is easily adapted to most SQL Server DBMS. Clearinghouses can also be exported to XML to allow the easy implementation of standalone systems.

[7] ASCII text files, dBase, ESRI shapefiles, JPG images, MrSID image server, common SQL DBMS servers, OGC Web Map server.

[8] TMWin also connects directly to datasets on a local network, whether registered with the clearinghouse or not, and caches Internet datasets to a local drive for offline use.

[9] TMWin kiosk mode is triggered by a command line parameter or keyboard combination, and offers a substantially simplified interface protecting the underlying data from modification by the user.

[10] Separate templates can be defined for each layer to override the MapSpace default, which in turn overrides the program default

References

Bahn, P. (ed.) 2000 *The Atlas of World Archaeology.* London: Time-Life Books

Dreyfus, H.L. 2001 *On the Internet.* London: Routledge

Egenhofer, M.J. 2002 Toward the Semantic Geospatial Web. *Proceedings of the Tenth ACM International Symposium on Advances in Geographic Information Systems,* pp 1-4 New York: ACM Press.

Jessee, C. 2001 Flash GIS: delivering geographic information on the Internet. Paper delivered at Digital Resources for Research in the Humanities. Computing Arts 2001, University of Sydney.

Johnson, I. 1999 Mapping the fourth dimension: the TimeMap project. In Dingwall, L., Exon, S., Gaffney, V., Laflin, S. & van Leusen, M. (eds) *Archaeology in the Age of the Internet.* British Archaeological Reports International Series 750. 21pp, CD-ROM.

Johnson, I. 2002 Contextualising Historical Information Through Time-Enabled Maps. *Internet Archaeology* 12.

Johnson, I. 2002 From metadata to animation: Web-based searching and mapping of cultural heritage information. *Proceedings of UNESCO World Heritage Centre 30th Anniversary Virtual Congress.* http://169.229.138.138/unesco/papers/3112_991-2002_10_unesco_alexandria.doc

Johnson, I. & Wilson, A. 2002 The TimeMap Kiosk: Delivering Historical Images in a Spatio-Temporal Context. In Arvidsson, J. & Burenhult, G. (eds) *Archaeological Informatics - Pushing the Envelope* British Archaeological Reports International Series 1016. Oxford: Archaeopress.

Kraak, M-J. & Brown, A. (eds) 2001 Web Cartography: developments and prospects. London: Taylor & Francis.

Lancaster, L.R. & Bodenhammer, D.J. 2002 The Electronic Cultural Atlas Initiative and the North American Religion Atlas. in Knowles, A.K. (ed) *Past Time, Past Place: GIS for History.* Redlands: ESRI Press. 163-177.

Monmonier, M.S. 1996 *How to lie with maps* 2nd ed. Chicago : University of Chicago Press.

O'Brien, P. (ed) 2002 *Concise Atlas of World History.* Oxford: Oxford University Press. 312pp

Software Tools for Indigenous Knowledge Management

Jane Hunter, DSTC Pty Ltd; Bevan Koopman, University of Queensland, Australia; and Jane Sledge, Smithsonian National Museum of the American Indian, USA

Abstract

Indigenous communities are beginning to realize the potential benefits digital technologies can offer with regard to the documentation and preservation of their histories and cultures. However, they are also coming to understand the opportunities for knowledge misuse and misappropriation of their knowledge which may accompany digitization. In this paper we describe a set of open source software tools designed to enable indigenous communities to protect unique cultural knowledge and materials preserved through digitization. The software tools described here enable authorized members of communities to define and control the rights, accessibility and reuse of their digital resources; uphold traditional laws pertaining to secret/sacred knowledge or objects; prevent the misuse of indigenous heritage in culturally inappropriate or insensitive ways; ensure proper attribution to the traditional owners; and enable indigenous communities to describe their resources in their own words. Hopefully the deployment of such tools will contribute to the self-determination and empowerment of indigenous communities through the revitalization of their cultures and knowledge which colonization, western laws, western cultures and globalization have eroded.

Keywords: Indigenous Knowledge, Cultural Rights, Rights Management Software

Introduction

Using multimedia technologies, indigenous groups have been able to record and preserve significant aspects of their cultures including languages (First Voices), ceremonies, dances, songs, stories, symbols, design, artwork, tools, costumes, historical photographs, film, videos and audio tapes (NMAI, Barani, Diwurruwurru). Documentation of indigenous knowledge and history has become an extremely important tool to ensure the survival and self-sustainability of indigenous tribes and cultures, and to provide evidence of past injustices and to support claims of original ownership.

Although digitization is ideal for sharing, exchanging, educating and preserving indigenous cultures, it also creates ample opportunities for illicit access to and misuse of traditional knowledge. It is essential that traditional owners be able to define and control the rights and access to their resources, in order to uphold traditional laws; prevent the misuse of indigenous heritage in culturally inappropriate or insensitive ways; and receive proper compensation for their cultural and intellectual property. Finally, it is essential that indigenous communities be able to describe and contextualize their culturally and historically significant collections in their own words and from their own perspectives.

In this paper we have investigated the application of IT security mechanisms to the rights management of indigenous collections. In consultation with staff from the NMAI Cultural Resources Centre (National Museum), we have developed a set of low cost, simple-to-use and robust software tools designed to enable the description, annotation and rights management of collections of mixed-media digital and physical objects belonging to indigenous communities. We have also developed a search, retrieval and presentation interface which retrieves different result sets, depending on the user's profile, and aggregates the results automatically into coherent multimedia SMIL (Synchronized) presentations.

Because of the enormous diversity of indigenous cultures, the system has been designed so that it can easily be customized to support the unique requirements of specific communities. The immediate future involves working closely with a small number of indigenous communities to refine the software to suit their specific requirements and projects. Eventually we hope to make the software downloadable from the Internet and freely available to indigenous communities for non-profit use and to provide training in its use and maintenance.

Background and Requirements

The work described here began with an investigation of the ability of current and evolving information technology tools for rights management (e.g., frameworks, markup languages, metadata models and standards) to support the unique needs of indigenous communities.We found that initiatives such as MPEG-21 (Multimedia Delivery Framework) (MPEG-21) and XrML (XrML) are primarily concerned with e-commerce and protecting the commercial rights of content owners.They are built on the premise of modern intellectual property law regimes and the notions of individual property ownership for a limited duration, ideas which are alien and detrimental to indigenous cultures. MPEG-21 and XrML do not support the specific requirements expressed by indigenous communities who need to protect indigenous knowledge or enforce tribal customary laws.

In an earlier paper (Hunter), specific extensions to XrML, in the form of customary constraints were developed to support the description of customary or traditional laws which commonly affect access.An analysis of tribal laws across Aboriginal and Native American communities revealed the following common factors or variables which may determine access to traditional knowledge:

- native/non-native restrictions;
- the user's membership in a particular clan or tribe;
- the user's status within the tribe (e.g., elder, initiate, child);
- the user's role within the tribe (e.g., dancer, artist, midwife, healer);
- the user's gender (male, female); o moon (menstrual cycle) restrictions; o pregnancy restrictions;
- the relationship of the viewer/user to the people, animals or objects depicted in the resource;
- the death of people recorded in a resource;
- human remains access restrictions - no access should be provided to images of human remains, or specifics about the disposition of human remains repatriated to tribal communities;
- the context in which the resource will be reused or reproduced.

Staff at the NMAI CRC also expressed the need to be able to document traditional care information associated with culturally significant physical artefacts in their custody. Kristina Dunman has described the meaning and importance of traditional care of American Indian artefacts in (Dunman). Jim Pepper Henry (Repatriation Program Manager at the NMAI) also provided the following list of traditional care constraints frequently requested by American Indian tribal representatives regarding the storage and handling of objects:

- directional orientation;
- segregation from other objects or other tribes' materials;
- storage on higher shelves, use of wooden shelving;
- cover from view with cloth or muslin;
- ceremonial feeding/dusting with corn pollen and/or other materials;
- cleansing with smoke (cedar, sage, tobacco, sweetgrass etc.);
- freezing restriction;
- low oxygen restriction (no CO_2 bubble or covering with plastic or other non permeable materials);
- sacred/ceremonial bundle dis-assemblage restriction (bundle to remain intact and unopened);
- ceremonial pipe assemblage restriction (bowl and stem should be housed disconnected).

Similar traditional care recommendations also apply to culturally sensitive artefacts belonging to Aboriginal and Torres Strait Islander communitiesbut held in museums, archives, cultural centres and keeping places.

The software should also enable the traditional owners to describe, contextualize and annotate resources in their own words, their own languages and from their own perspectives. The importance and value associated with enabling spoken annotations (in addition to textual annotations) was multi-faceted:

- Spoken annotation tools reinforce and support the oral tradition which is so strong in many indigenous cultures;

- Spoken annotations are an easier and more natural interface for user input than keyboards, particularly for communities with low computer literacy and poor keyboard skills;

- Spoken annotations represent new language resources which can be used to help preserve threatened languages;

- Photos and videos can act as a trigger for the Indigenous elders to record their stories as spoken annotations to the visual resources.

In addition, users should be able to view or listen to associated annotations which are clearly attributed to individuals. This approach supports the unambiguous documentation of all views/perspectives - even if they are different or contradictory. The software explicitly displays "who said what and when" rather than displaying only the view of a Museum Curator which may have been deduced from a number of different sources of varying reliability.

Figure 1 illustrates how we envisage the software being used to assist with the protection, preservation and repatriation of indigenous knowledge and artefacts which are being shared, exchanged or returned from museums, archives, private anthropologists' collections and cultural institutions back to their original owners. (N.B. We are not considering the repatriation of human remains within the scope of this project.)

No single approach is applicable to the repatriation of indigenous information, knowledge and/or artefacts. For example, Australian law differs from US law, and additionally, each tribal community will have its own unique needs and requests. In the United States, while the Native American Graves Protection and Repatriation Act (NAGPRA) specifies the types of objects and sites to be protected and/or repatriated, it is expected that some tribal communities will want access to the records of all objects in museum collections associated with their community and be satisfied with digital surrogates and access to physical objects when requested. Going beyond the requirements of NAGPRA, the NMAI has established a Culturally Sensitive Collections Care Program to respond to areas of concern of Native peoples with regard to the maintenance, presentation, and disposition of sensitive materials and information in the collections of the museum. Jim Pepper Henry, Assistant Director for Community Services at NMAI says:

> This Program is to be implemented with regard to the wishes and concerns of indigenous communities and traditional leaders and structured within the boundaries of the obvious and reasonable limitations of the institution. The basis for this program is formed through consultations with official tribal representatives, tribal elders and traditional leaders, and museum staff with respect to museum policies and procedures, and US federal legislation including the National Museum of the American Indian Act, American Indian Religious Freedom Act (AIRFA), Native American Graves Protection and Repatriation Act (NAGPRA), and the Collections Policy of the National Museum of the American Indian. The Program Committee is cognizant of the fact that it is operating within the constraints of a national museum, and acknowledges that in some instances, dependant upon the beliefs and feelings of the Native community concerned, the proposed implementation of sensitive collections care practices within the museum setting may be deemed inappropriate or presumptuous. In such circumstances, the Committee will endeavour to abide by any alternative direction offered by the concerned Native group.

The aim of the software described here is to support such a Program.

The envisaged usage and application of the software as described above, introduced certain requirements and design constraints which needed to be satisfied. These included:

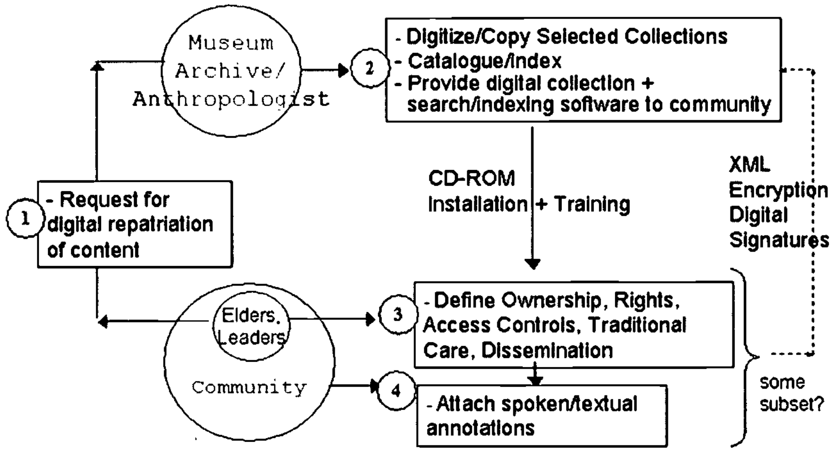

Figure 1 – Envisaged Software Usage Scenario

- Security mechanisms - because of the sacred/secret nature of the content with which we are dealing, it is essential that the IT security mechanisms which are employed be impenetrable and reliable;

- Simple user interfaces - many of the potential users of this system will have low computer literacy, so simple intuitive user-friendly interfaces are essential;

- Robustness - the system must be able to stand up to the rigours of unexpected input by users with little prior computing experience;

- Low cost - in order to make the software open source and accessible to indigenous and grassroots communities, it must be built as inexpensively as possible, using tools which are ideally free;

- Interoperability - the software tools should be built on international standards - Dublin Core (Dublin), CIDOC CRM (CIDOC), MPEG-21 (MGEP-21), XrML (XrML)- in order to ensure maximum interoperability between disparate databases;

- Portability - it should be able to run on a range of platforms and operating systems. Java (JDBC, JSP), XML and SMIL have been used as the software development environment to ensure transparent portability across platforms;

- Flexibility - The customary laws and intellectual property needs of traditional knowledge holders vary enormously among indigenous communities throughout the world. Quite often the views within a single clan can vary significantly, and they may also vary over time. Our system attempts to support the common notions associated with traditional laws within Indigenous communities. In addition, we have provided Schema editing tools in order to provide maximum flexibility and to enable easy customization of the software.

- Scalability - the size of indigenous collections (particularly within cultural institutions) can reach hundreds of thousands of items. The software should be capable of efficiently enabling metadata/constraints to be applied across large sets of resources, individual resources or regions/segments within resources for either individual users or user-groups.

Although a number of other projects have investigated the application of information technology tools to the protection and management of indigenous collections according to customary laws (Ara, Sullivan, Digital, Special), they have not approached the problem using international metadata standards nor developed generic yet flexible systems which are capable of supporting indigenous communities globally, but easily customized to support the particular local customs.

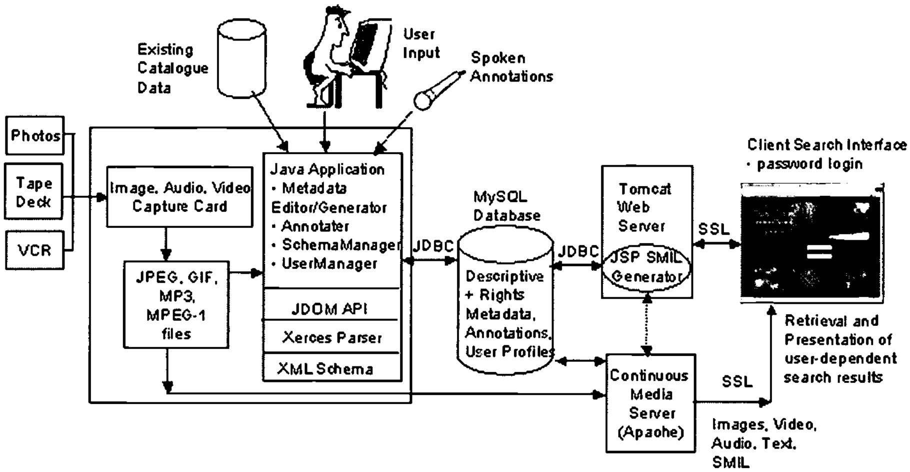

Figure 2 - System Architecture and Workflow

The remainder of this paper describes the software which has been developed to support the requirements specified above.

System Architecture and Components

This section describes the architecture of the software system, its separate components and the overall process/workflow described in Figure 1.

The system consists of three major components:

1. The Metadata Editor/Generator;

2. The Database;

3. The Search, Retrieval and Presentation Interface.

Figure 2 illustrates the interfaces to these components and the technologies used to build them and integrate them into a single coherent system.

The Metadata Editor/Generator

This component of the system enables users to input the descriptive, rights and tribal care metadata associated with the objects (either physical or digital) and to attach spoken or written annotations to specific objects.

Users require a login ID and password to run this software component and, depending on their privileges/user profile, may only be permitted access to certain functionality.

Figure 3 illustrates the three types of metadata which can be input and the users/groups who we envisage will provide them.

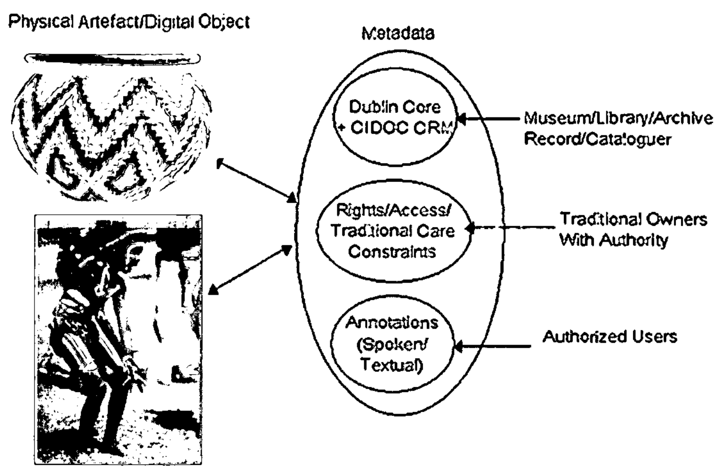

Figure 3 - Metadata Types and Sources

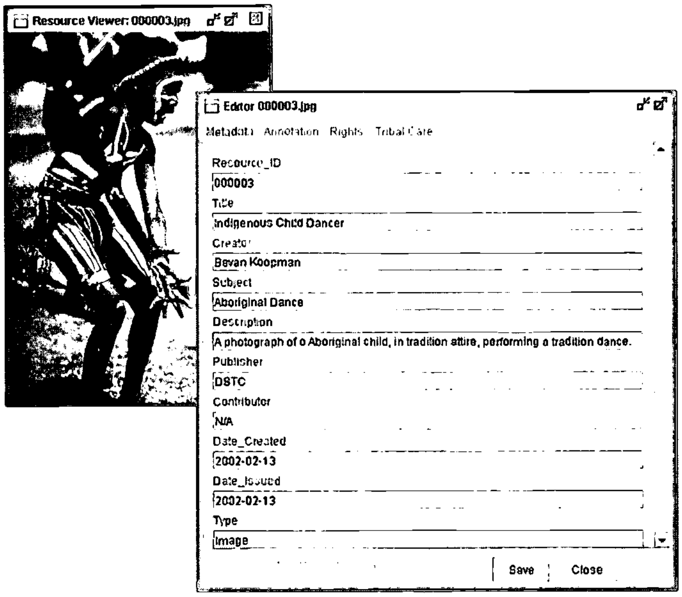

Figure 4 - Descriptive Metadata Input

The Descriptive Metadata consists of Dublin Core [Dublin] plus some additional optional elements from the CIDOC CRM [CIDOC] for describing physical museum objects, such as material, technique and dimensions. Figure 4 shows the user interface for inputting or editing descriptive metadata. Alternatively, some or all of the descriptive metadata could be acquired from existing museum/library/archive database cataloguing information.

The Rights Metadata will be provided by the traditional owners or elders of the indigenous community to whom the resource belongs. Only those users with the required access privileges will be able to input or edit the rights metadata. Support has been provided to enable the definition and application of restrictions based on:

- the user's membership of a particular clan or tribe;
- the user's status within the tribe;
- the user's role within the tribe;
- the user's gender;
- the relationship of the user to people, animals or objects depicted in the resource;
- the death of people recorded in a resource;
- the context in which the resource will be reused or reproduced.

Figures 5a, 5b, 5c, 5d, and 5e show the user interface for defining access restrictions which depend on the user's tribal affiliation, gender, role and status.

In addition to the descriptive and rights metadata, annotation tools which enable indigenous communities to describe resources in their own words have been provided. Authorized users can input, record and attach either textual or spoken annotations to specific objects or resources. The ability to enter spoken annotations is especially useful and valuable - it provides a natural user interface which supports the oral tradition and allows users to express their stories in their own words and languages. It also acts as a trigger for generating new knowledge and enhancing existing knowledge. Figure 6, below, illustrates the annotation interface. Users can also browse the list of clearly attributed annotations and view/listen to who said what and when about a particular resource.

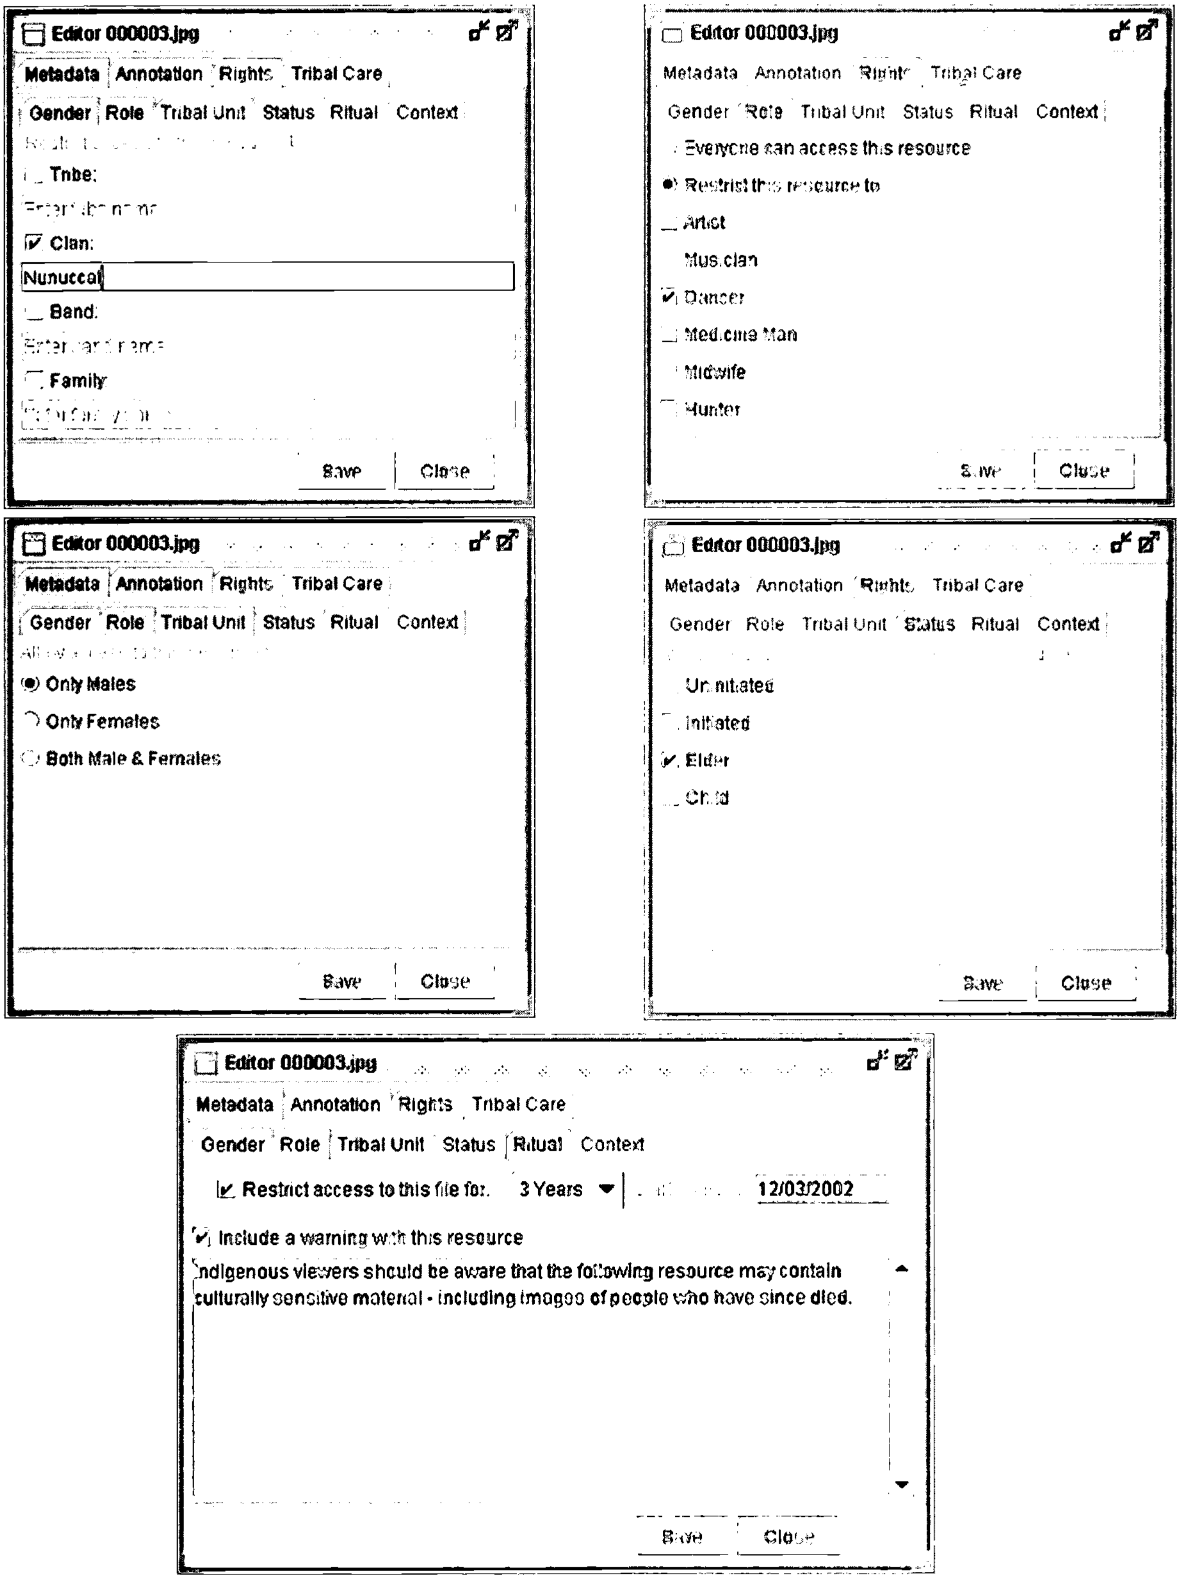

Figure 5 - User Interfaces for Restricting Access According to Customary Laws

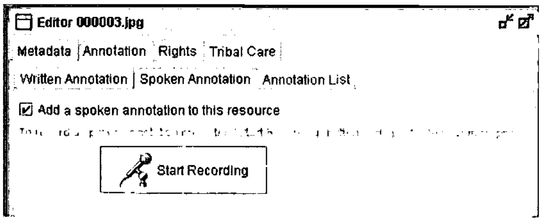

Figure 6 – Spoken Annotation Input

In response to a request from both museum staff and representatives of Indigenous communities, we have also added support for the specification of Tribal Care recommendations. The Tribal Care metadata tool enables Indigenous communities to define the ways in which culturally significant or sensitive physical artefacts should be treated - ways which are acceptable to the cultures represented. Elders are able to define particular spatial orientation, storage and preservation recommendations to ensure the respectful tribal care of physical artefacts. While the actual implementation of traditional care specifications may be difficult, many museums are attempting to integrate Indigenous beliefs and requests into museum practices, and our software will hopefully facilitate this process.

Where possible, the specified access restrictions will be enforced by matching them against the profile of the user. Every user of the system requires a login ID and password. When being allocated a user ID, users will also need to provide supporting documentation to prove their claims of tribal affiliation, status, etc.

The User Manager software component illustrated in Figure 8 allows the systems administrator to add or remove users and to edit their profiles. Authorized elders will be able to enter the rights constraints, and authorized users will be able to attach annotations. Hence the software affects not only what resources users can access and view but also what metadata tools they can access and hence what metadata they can enter.

Because of the enormous diversity of indigenous cultures, the system has been designed so that it can easily be customized to support the unique requirements of different communities. Customization is carried out through the SchemaManager tool. Community elders can add new constraints, or remove or refine existing constraints, depending on the traditional laws of their community. An XML Schema [XML] is saved to reflect their particular metadata requirements and rights constraints, and the user interface is then generated from the saved/selected schema. Figure 9 illustrates the Schema Manager user interface.

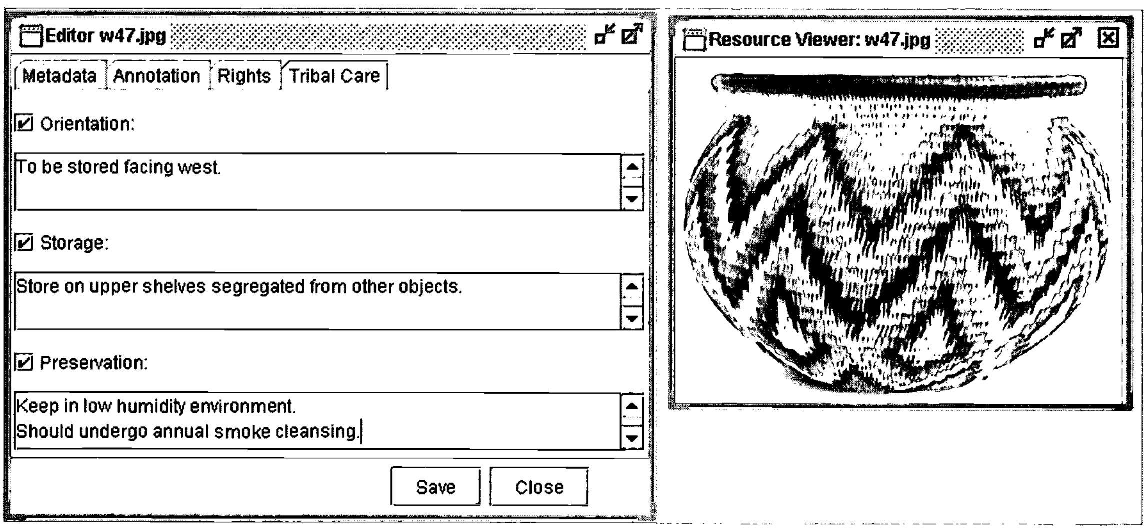

Figure 7 - Tribal Care Specifications

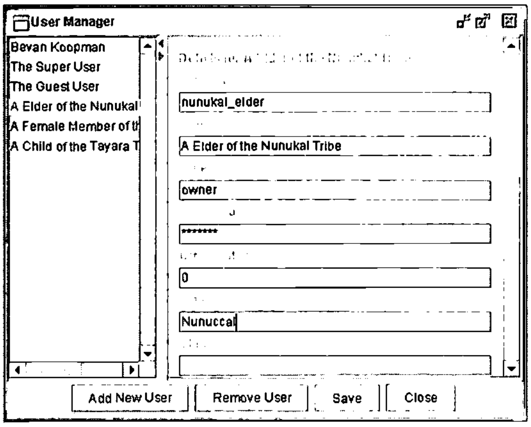

Figure 8 - User Manager

The Database

Saved metadata is stored in relational tables in a MySQL database which is connected through a JDBC (Java Database Connectivity) API [Java]. It is envisaged that, because of the sensitivity of the data, this database will not be accessible through the Internet. It will be stored on CD-ROM or the hard disk of a stand-alone supervised workstation, within a cultural centre, keeping place or supervised building (such as a library or a school) within the community.

In addition to the metadata which is explicitly saved to the database through the user interface, meta-metadata is also recorded - all changes to the metadata, who was responsible, and the date/time of the changes are recorded within the database. This represents an important component of the system's built-in security framework.

The Search, Retrieval and Presentation Interface

A search, browse and retrieval interface to the collection was built using standard Web Browser technologies (Internet Explorer, Netscape) for the user interface. The advantages of using standard Browsers for the search interface are their familiarity and widespread availability and the lack of re-engineering necessary should collections eventually be disseminated over wider networks. To access the collection, users must have been allocated a login ID and password and a user profile. The steps below illustrate the typical procedure which users would follow when searching and browsing an indigenous collection:

1. A user logs onto the system using a secure password. Associated with each user is an authenticated user profile which includes information such as tribal/western names, native/non-native heritage, tribal/clan membership, gender, status, role, etc.;

2. The user performs a search on a particular topic, e.g., dance; (See Figure 10);

3. The software then searches the title, subject, and description metadata associated with each object in the collection, for the specified search term (e.g., dance);

4. For those objects whose metadata matches the search term, the software compares the objects' rights constraints with the user's profile to determine whether or not the user is permitted to

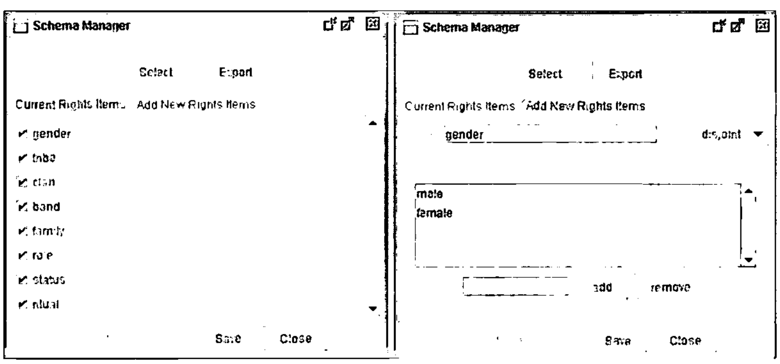

Figure 9 - Schema Manager

access this object. If so then this object will be added to the result set;

5. The list of results/objects which match the search term and which the user is permitted to access is then displayed - along with any rights constraints, which appear as icons (see Figure 11);

6. The user can click on individual objects to view/play the object and to view the metadata details and any annotations;

7. The users can select objects of particular interest and add them to a personal collection;

8. The software automatically aggregates those mixed-media objects selected by the user (images, audio clips, video clips, text), and dynamically generates a SMIL (Synchronized Multimedia Integration Language) [Simulated] presentation which is delivered to the user (see Figure 12);

Future Work and Conclusions

In this paper we have described a software system developed as a result of consultation with representatives from Indigenous communities and staff from museums, archives, libraries and cultural centres in both Australia and North America. However at this stage the software remains relatively untested within real world applications or real communities. Hence the immediate future involves working closely with a small number of indigenous communities to determine:

- whether the software can satisfy the unique requirements of particular Indigenous communities for the management of their culturally sensitive collections or whether further extensions/refinements/modifications will be required;

- whether the dynamic political, social and trust issues (e.g., lack of agreement on access rules, vali-

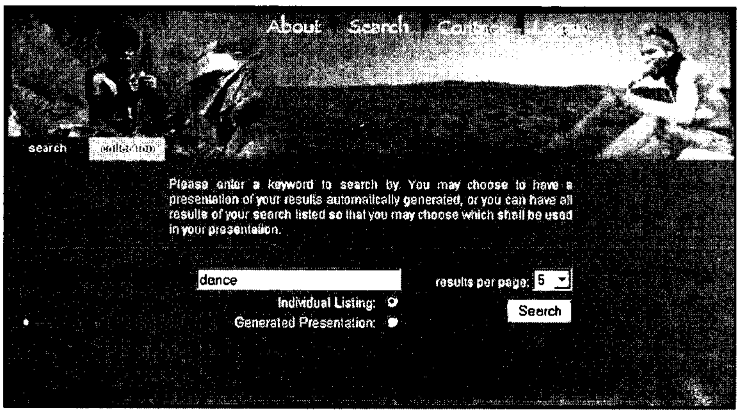

Figure 10 - Search Interface

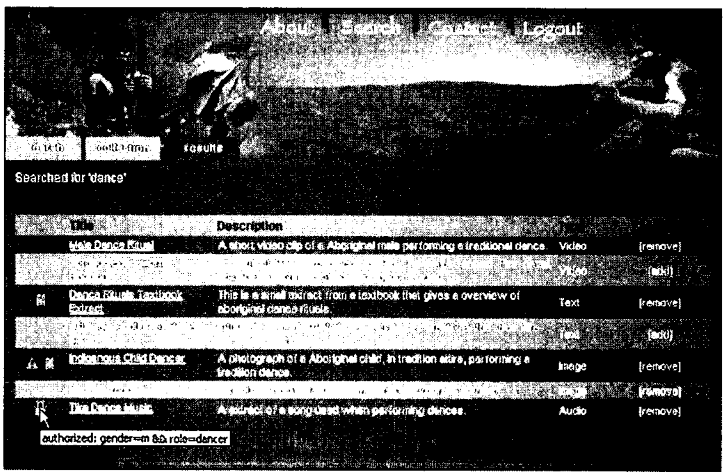

Figure 11 - Search Results

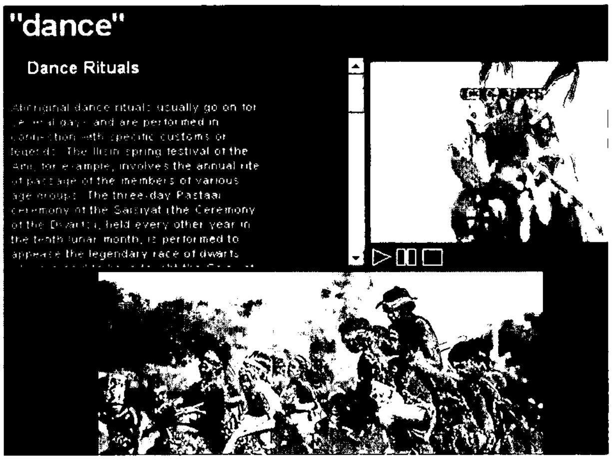

Figure 12 - Multimedia Presentation of Results

dation of individual claims of authority, authorizations within a tribe or community) are greater than the technical problems associated with enforcing them;

- the proper procedures required to enable the successful and beneficial application of the software to the preservation, description, protection and annotation of indigenous cultural collections by the traditional owners;

- the types of scenarios, situations, collections and communities to which the software is most suited and of maximum benefit;

- whether additional security mechanisms such as XML Encryption [XML Encryption], XML Digital Signatures [XML Digital], SAML (Security Assertion Markup Language) [Security], SSL (Secure Sockets Layer) [Secure Socket] and watermarking techniques, would be applicable and could be trusted to ensure secure access to and transfer of sacred/secret data over networks between distributed remote locations and a common centralized repository.

Given a positive response from the community trials, we plan to make the software freely available and downloadable from the Internet for non-profit use by indigenous communities and to provide training in the usage, refinement and maintenance of the system to interested groups.

Acknowledgements

The work described in this paper has been carried out as part of a Queensland Smithsonian fellowship, funded by the Queensland Government and a collaboration between the Distributed Systems Technology Centre and the Smithsonian's National Museum of the American Indian Cultural Resources Centre (NMAI CRC). Thanks also to Angelina Russo and students from UQ's Information Environments program for their Web search interface design.

References

Ara Irititja Archival Project, South Australia, http://waru.org/arairititja

Barani - Indigenous History of Sydney City, <http://www.cityofsydney.nsw.gov.au/barani/main.html>

CIDOC Conceptual Reference Model, <http://cidoc.ics.forth.gr>

Digital Collectives in Indigenous Cultures and Communities Meeting, Hilo, Hawaii, August 10-12, 2001, <http://si.umich.edu/pep/dc/meeting/meeting.htm>

Diwurruwurru - "Our Message Stick to the World", <http://arts.deakin.edu.au/Diwurruwurru>

Dublin Core Metadata Initiative, <http://www.dublincore.org>

First Voices Project, <http://www.firstvoices.com>

Hunter, J. (2002). Rights Markup Extensions for the Protection of Indigenous Knowledge, 11th International World Wide Web Conference, Honolulu, Hawaii, May 2002, <http://archive.dstc.edu.au/IRM_project/paper.pdf>

Java Database Connectivity(JDBC), <http://java.sun.com/products/jdbc/>

MPEG-21 Overview, <http://mpeg.telecomitalialab.com/standards/mpeg-21/mpeg-21.htm>

National Museum of the American Indian(NMAI), <http://www.nmai.si.edu/>

NMAI Exhibits online, <http://www.conexus.si.edu>

Security Assertions Markup Language(SAML), <http://www.saml.org/>

Secure Socket Layer (SSL), <http://wp.netscape.com/eng/ssl3/>

Special Issue on Digital Technology and Indigenous Communities, D-Lib Magazine, March 2002, <http://www.dlib.org/dlib/march02/03contents.html>

Sullivan, R. (2002). Indigenous Cultural and Intellectual Property Rights - A Digital Library Context, D-Lib Magazine May 2002 Volume 8 Number 5

Synchronized Multimedia Integration Language (SMIL), W3C Recommendation, 07 August 2001, <http://www.w3.org/TR/smil20/>

XML Digital Syntax and Processing, W3C Recommendation, 14 February 2002, <http://www.w3c.org/xmldsig-core>

XML Schema Language, W3C Recommendation, 02 May 2001, <http://www.w3.org/XML/Schema>

XrML, eXtensible rights Markup Language, <http://www.xrml.org>

Dublin Core: The Base for an Indigenous Culture Environment?

Liddy Nevile, La Trobe University, and Sophie Lissonnet, James Cook University, Australia

Abstract

A day in Cape York, in the far north east of Australia, can change the life of a modern Australian. In that time, one can see hundreds of examples of rock art that are up to 36,000 years old, sharply contrasting the history of Indigenous people and the immigration of Europeans. One such visit led to a proposed collaboration between the Quinkan Culture Elders and a team of metadata researchers. The researchers proposed a Qualified Dublin Core style catalogue to be used to identify and record examples of Quinkan Culture so Elders could at last gain access to information needed to manage the proliferation of unauthorised publications about Quinkan culture, and to 'bring back home' cultural representations. In addition, the catalogue would allow the Elders to make decisions about publishing their own representations. This paper describes the journey of members of the team developing 'Matchbox', a cataloguing system, as they have sought a way of using Qualified DC metadata (QDC) to describe, collect, and represent Quinkan Culture. Of interest in this paper is how developing a QDC representation has led to questions of cultural definition and, simultaneously, of the use of technologies such as HTML, XML and RDF.

Keywords: Qualified Dublin Core, HTML, XML, RDF, Quinkan Rock Art, culture

Introduction

A day in Cape York, in the far north east of Australia, can change the life of a modern Australian. In that time, one can see hundreds of examples of rock art that are up to 36,000 years old, sharply contrasting the history of Indigenous people and the immigration of Europeans.

One such visit led to a proposal for collaboration between the Quinkan Culture Elders and a team of metadata researchers[1]. The researchers proposed a Qualified Dublin Core style[2] catalogue to be used to identify and record examples of Quinkan Culture so Elders could at last gain access to information needed to manage the proliferation of unauthorised publications about Quinkan culture, and to 'bring back home' cultural representations. In addition, the catalogue would allow the Elders to make decisions about publishing their own representations.

In the nearest city (200 miles away), Elders attended a computer workshop and accessed the Web. Most exciting was the viewing of a recent video of some Elders talking about a honey tree. They enjoyed sharing their stories of collecting honey. They decided a multimedia catalogue would make it easy to prompt new stories and explanations without the need to travel far out into the country.

Meanwhile, working with the Dublin Core Metadata Element Set (DC) has been found by the authoritative museum community to lack much of the richness that is valued by curators. CIMI[3] several years ago energetically investigated the possibilities and found it necessary to modify and extend DC. This work was done with early versions of DC, however. Recently, DC has gained from experience and changed so that it may, indeed, now perform more usefully in this role. The Quinkan researchers are investigating this.

This paper describes a journey undertaken in a university office by some researchers through a myriad of disciplines as they have sought, through literature and so far a very few interactions with Indigenous Elders, a way of using Qualified DC metadata (QDC) to describe, collect, and represent Quinkan Culture.

Of interest in this paper is how developing a QDC representation has led to questions of cultural definition. This is not a paper about how Indigenous people have worked on their cultural interests. It is about how the questions being asked by the researchers have changed. In fact, these questions appear as indicators of the change in their understanding. Why such changes are reported here is because they seem to offer a way of thinking about the use of metadata-based technologies to explore the nature of tools that may be useful to those working in cultural contexts.

Context

We held a series of meetings and exploratory visits to the far north east of Australia, where the communities are small and Indigenous people have been involuntarily scattered. We, a small group of European Australians, proposed a project that we hoped would serve as a useful resource for our Indigenous colleagues and help them repatriate their heritage. We are not experts in such matters, with the exception of one of us who is a specialist in the archaeology of the region.

To us novices, it seemed that a catalogue of cultural heritage would make it easier for the Indigenous people to control their heritage. There was a rewarding workshop in which Indigenous Elders showed their pleasure in sharing interpretations of events and objects, and we assumed from this that a catalogue might be a good organising tool and stimulus for gathering such stories.

Naively, perhaps, we thought Quinkan culture would somehow be amenable to being catalogued. It seemed that all that would be required was something like an extension of an archaeological catalogue.

Sadly, there no longer are rock artists, and relevant cultural practices are scarce. Today, Elders tell stories of hunting, of gathering bush tucker, of lives in which they have unfortunately been victims of the invasion of their country. They gather to celebrate their dance culture. They take people to see some of the many Rock Art paintings. There has been carbon dating of some paintings, scores of anthropological, geological, and many other forms of study of the region, and there are data in many forms and locations world-wide.

The government is building an 'interpretive centre' in the region, but what is to happen there is unclear. This 'open' museum may be a collection of tangible and intangible assets, a presentation site, a meeting place. Whatever, it is in its infancy as a cultural institution. One idea is that the Quinkan catalogue will be a living exhibit, an environment, in this small space, and possibly a cultural heritage management tool.

The First Question

Given that the most obvious 'cultural assets' in the region are Rock Art paintings and petroglyphs, our focus naturally gravitated towards these. This is not all there is in the region, of course, but such a vast number in apparently original condition is attractive as a 'collection', a set of objects that can be identified, described, possibly even quantified. It is also interesting to note that these paintings are fixed; their location is stable and discernible in great detail using modern geo-spatial tools and techniques. At first glance, it might seem that the first question would be how to describe the location of the paintings, and then everything would follow.

Quinkan Rock Art occurs in what have been called galleries, usually open shelters high up on steep overhanging rock outcrops. The areas are appealing because of their aspect, catching the breezes and providing shelter, and frequently also having good views of the surrounding country. Galleries contain as many as hundreds of paintings, often superimposed on one another, with remarkably fresh-looking paintings in abundance. Fortunately, although there are almost no resources available to protect this heritage, it is isolated and has not suffered, as it might, from extensive tourism. Like the flora and fauna of Australia, the Rock Art sites are geologically frail, and human visitations cause significant damage to the sites themselves, as well as to the paintings.

Not all paintings or galleries have been identified and had their location recorded. Some have had minute details recorded about them, while others are known to be sort of 'over there'. But the idea of a stable identification system with descriptions of material culture is attractive, and modern technology offers great precision to the task of recording relevant locations.

So using location with sufficient detail to develop a geo-spatial topology could solve our problems. This is, after all, a typical approach: every resource in our collection could have the location of its content so well specified that it would only need to be differentiated from others that are related to ex-

actly the same location. Trouble is, Elders do not want their paintings, or presumably their culture in general, identified in this way. The way we want to describe locations is, a matter of cultural preference, and geo-spatial science is not a traditional way. In Australian Indigenous cultures, land is not just a commodity, and location is not just a number. The sites where paintings are located are the same places where sacred objects and spirits are likely to be found. People and their country, locations as others think of them, are inextricable intertwined. Access is often limited to a few people, and the rules for access are not published.

For non-Indigenous researchers who work on Indigenous cultural heritage in the area, recording and revealing locations has no more interest than for their Indigenous colleagues. The paintings are in remote country, perhaps several days' walking time from the nearest accessible track, and that might be several days' driving from the nearest town. Small helicopters can sometimes land directly near sites. But site visitation is always at a cost, and sadly, it is the condition of the paintings that suffers most when dust is disturbed by visitors and abrasion wears out yet more of the precious images, or when ancient meeting places are disturbed.

What if we were dealing with a Rock Art painting, a chart, a photograph, a scanned image, etc and all were at exactly the same place? The geo-spatial information would not be sufficient to distinguish among the objects just listed. The differences among the objects would be of form, perhaps ownership, possibly content as viewed by different people, and more. Recording these differences, or identifiers, is the primary work of the Dublin Core Metadata Element Set (DCMES). Every object described by DCMES has a unique resource identifier (a URI) and a range of other identifiers.

Seeking Answers to the First Question

So the first question was: how should the objects to be catalogued be described using a DC-based system?

Starting in what seemed the obvious place, the team consulted the resident archaeologist who, as an archaeologist, has focused on particular examples of material culture, paintings.

Typical archaeological work describes in detail some aspects of the paintings and has at times depended upon descriptive vocabularies that allow for classification of aspects of paintings. A good example is the use of colour classifications that allow archaeologists to make comparisons of the colouring of paintings. Given that all the paintings are in the open air, the time of day makes a significant difference to their appearance. Observations of colour need to be time-stamped and seasonally-marked as part of their codification if they are to be useful. The development of useful colour classifications has therefore been of interest to archaeologists, and there are specialised and idiosyncratic classification schemata.

So to classify all the Rock Art paintings by colour should be a relatively objective activity, and there are established criteria. The same is true of other aspects of the paintings: the date, the location, the shapes of the objects within the paintings, even photographs of the paintings, drawings of them, stories about them, everything could be described. Not all the items would be paintings, of course. In fact, we anticipate stories about paintings and stories about those stories, something like annotations of interpretations or descriptions. But still these objects can be described. Each one is a unique object and with a good description, can be identified.

The information scientist in the team is versed in metadata and had no trouble building a metadata profile that provided for elements and sub-elements, with controlled vocabularies, specialised formats, etc. This work made good sense and could be neatly represented in a hypertext Web page with pull-down menus and active buttons.

The 'describing process' was reduced to the filling in of a form. The description would be more or less useful depending upon how well the elements and qualifications of them were chosen. The metadata application profile (MAP) for each item in the catalogue had be developed in the light of experience over many years, even centuries, on the part of museums, libraries, scholars and others.

Describing objects in this way has been, for a long 'Web' time, the primary use of metadata, particularly of Dublin Core (DC) metadata. Finding the criteria that make it possible to distinguish an object from all other objects with which it might be confused is a primary role for a metadata profile. If a system can use the description of each item in an appropriate way, it will have a handle by which it can perform functions and provide services, such as discovery.

A good example is provided by a resource that is a photograph. If the description is complete, it will be clear not only who and what are in the photograph, the 'aboutness' of the photograph. It will also be discernible which version of the photograph it is, where it is, and so on, and the photograph will be distinguishable from another that looks similar but is, in fact, a second print from the same negative.

We have come to think of this as a first dimensional activity. Unless an item can be uniquely identified, nothing more is possible, so its description is the crucial first requirement, its DNA. The description that we produce for items in Matchbox can be represented completely in Hypertext Markup Language (HTML[4]). It is capable of being written in plain text with tags to show what each bit means, and it can easily be read and understood by humans. Significantly, the tags can be read by machines so they too can tell what bit goes where. The meaning of what is between the tags, the content that is added to the MAP, is, of course, opaque to the machines.

The First Question Leads to the First Dimension

So working on the first question leads to a set of textual statements that are meaningful to humans and computers and can isolate a chosen object from others.

First dimension mapping (filling in of a MAP) is primarily for discovery. It can also be used to do simple things like pointing to links between objects. In other words, first dimension maps can be used in a first dimension way, for discovery, because they can be used for matching and thus for filtering, and so they can support discovery.

An object might have the special DC tag that contains information about how one object relates to another. Suppose this digitised photograph is of a painting that records a visitor in a gallery that contains a Rock Art painting. The digitised photograph can be linked syntactically (by a machine) to the photograph of the painting if both objects are registered in the same catalogue: the digitised photograph will have the relationship 'is format of' to the photograph. Significantly, the relationship is between one object and another; for instance, a digital file and the photograph that was scanned. There is no machine-readable mechanism for recording the semantic link between the photographs. That is, that the two photographs and the chart are all of the same Rock Art painting is not machine accessible information. The relationship of the semantics, or content, is not expressed in HTML although by filtering the HTML an application could get access to information to make this link. (The important thing to note is that HTML does not convey the link. It is not a part of the description of the photograph, but rather a connection that has to be made somehow.)

It is critical that the object's description contain the necessary information and that the information be in good form and understandable, with ways of deciphering it where necessary, such as dependence on a documented controlled vocabulary. It is not critical, however, that the information be in a simple form such as the proposed HTML version (referred to above). It may be in a more complex form, expressed in eXtensible Markup Language (XML[5]) or even using the Resource Description Framework (RDF[6]). Where this is the case, there is generally also a human-readable version made available by the system responsible for the representation because, although the text content is within the XML or RDF, it is surrounded by a lot of other syntactical clutter that may make reading it difficult. The form of expression is not important for its first dimensional use, for discovery that is based on unique identification, for example. Whether the additional information contained in the XML or RDF format is a problem or not depends upon the system using it.

The Second Question

The second question relates to how the semantics of the descriptions can be made to carry information about the structure of the descriptions. How can it be made known to a machine that this object is a scanned version of that object? It turns out that finding a way to represent structure is the next problem.

If I am looking for something and do not know its DNA, or MAP values, I may want to 'browse' through what is available until I find something that will do for me. This is the typical position of many users of Web navigation; they know they want something but either have not spent the time or do not know how to frame a discovery question that will guide them to a satisfactory object (containing the required information). This user is doing something quite different from those who know exactly what objects they are s looking for – the user in this case is usually looking for information about something but not sure what form that information will take (and maybe does not care anyway). The mapping that will provide this user with what is wanted will have to be used somehow to link objects in some hierarchy that will lead the user towards material satisfactory for the purpose. In fact, often the user likes to vary the discovery process and tries at times to narrow the field with a search during the browse, but inasmuch as he or she is using a browse navigation system, the object's description MAP will be essential. It will be used to show that this is a scanned version of that, for instance, because the two objects will be identified as having a relationship to different parts of a single hierarchy of production of digitised images.

Question Two Works to the Second Dimension

In the Matchbox case, it is not clear how users will want to browse, or, more correctly perhaps, what they will expect to find as the underlying structure of a browsing activity. Users do not always see the structure, but it is important that they learn to 'drive' it, and this usually comes from the effective development of a structural representation in the users' minds. Users do not like suddenly being jumped to a resource or set of resources that do not have what appears to them as a logical connection with what they thought they were asking for, or the resources they thought they were already viewing.

Some Matchbox users do act in predictable ways: those who work within established disciplines usually have learned to respond to the disciplinary taxonomies and terminology and can make fairly accurate guesses about what will happen if they are browsing up or down a hierarchy of criteria. It is also fairly predictable what type of criteria they will use: title, format, aboutness, authorship, etc., in many cases.

The Dublin Core Metadata Application Profile captures this information well. Once there is more specificity, qualifications to the DCMES can focus the mind appropriately, especially if there are additional elements that are useful in the context (for instance, the CIMI profile for museums). But in order to be machine readable, this structure must be declared in syntax that makes it accessible to machines. So the team has to express its qualified, extended Dublin Core-based metadata application profile for Matchbox in a way that will convey the relationships. XML, as recommended for the DCMES, makes it possible to declare both the relationships and the data.

The photograph of the person referred to above might contain information such as that the person 'is the child of xx' and 'is the parent of yy'. This is interesting but not readable information if it is expressed in HTML: the semantics are opaque. But if the semantics are expressed in a machine-readable way, the photograph, upon inclusion in the collection, can be linked by a set structure such as a family tree to pictures of the parent and child if they are in the collection.

Ideally, where elements are found useful and derived from established metadata application profiles, it should be easy and most accurate to simply import them into a new, combined, in this case Quinkan, profile. eXtensible Markup Language (XML) makes this possible, so instead of recording the maps of resources in HTML, it becomes necessary to record them in XML so their structures will also be recorded and thus decipherable.

XML is such a versatile language that many people can be using compliant, valid XML, and recording the structure of their information about their resources, and not be constrained by having to use the same way of representing their information. If MAPs from different sources are to be combined, however, there will be a problem with also conveying the structure that is associated with those MAPs. XML does not provide an easy way to solve this as the structure needs to be declared in advance, and it may be broken by the process of combining bits and pieces. In other words, the structure and the information are so closely coupled in XML that there is a risk the coupling will be broken when the XML is fragmented.

In the Matchbox case, this means recombining the pieces of others' MAPs with the new, special Quinkan pieces, to work in an integrated structure to be provided by the team. This structure, or taxonomy as we refer to it, will determine the browse structure. In fact, we may provide more than one to suit more than one group of predictable users. Some taxonomies, particularly those that are established elsewhere as useful to disciplined users (archaeologists, anthropologists, etc.), have already been identified and will be used to organise the semantic content of the resource maps, but they will need to be augmented by more general user-oriented taxonomies.

Just as HTML was closely related to the opening up of the first dimension of discovery, eXtensible Markup language that provides structure has made possible the second level. Without a structured description, there is not the control required. But once in XML, it does not matter if there is more structure than is necessary. For instance, if the markup complies with the Resource Description Framework (RDF), it is XML and useful at both the first and second level but may have added advantages.

Pre-determining the structure of the resources unfortunately does not work well in the Matchbox context. Our major challenge is not to find how to catalogue resources so that those who work within well-defined disciplines can find them, but to extend the opportunities for use of the resources and particularly to find ways of making them useful to the Traditional Owners of the Rock Art. We had no idea how such people would want to classify information about their heritage.

So we moved to question three.

Question Three

As soon as we asked how the resources should be organized, we found ourselves asking what they would be used for. If Matchbox is a cultural heritage tool, it should do things for its users. The problem is, we don't know what they will want to do.

Perhaps users will just want to find a resource or find if a resource exists. Perhaps they will want to know how it came to be in the form it is in. But what if they want to do with resources something like they do everyday – i.e., use resources in a cultural practice?

It is a well-known activity to try to find a link through friends of friends from one person to another: 'I know Mary and she knows Tom'. 'I am the daughter of Fred and he is the son of Mabel, so Mabel is my grandmother.' These are examples of the sort of relationships that people make between information, of cultural practices.

Australian Indigenous people often have long histories of their family in their memory, and they can relate these to people with whom they are interacting. They also have kinship structures that are different from the standard atomic family groupings that characterise modern western society. Linking people in photographs in some ways associated with indigenous kinship structures might prove to be of interest to Matchbox users in the future. So far, such structures are not defined for use in Matchbox, but they may be in the near future.

The Third Dimension

A simple example follows. It is not something that the Indigenous people have said they want to d,o but it is proposed as typical of something that might be of interest, so it can be presented to them.

Libby Miller and colleagues have developed Co-Depiction[7], an RDF system that takes a photograph, with the e-mail addresses of those in the picture, and uses this information to map from one person in that picture to a person in another picture by working on the friends-of-a-friend principle. If I am in a photo with Mary and Mary is in a photo with Jack, then two photos can be used to automatically link me to Jack.

In the Matchbox example, the idea is that as Indigenous people often use their genealogy to identify themselves, it might be interesting to use a yet further constrained friend of a friend principle to link photographs as they are placed in a collection. This could be achieved by having a structure, as described before, and having the new photograph deliberately located in relationship to that structure, each photo being described as containing an image of a person in a particular position in a family tree. But suppose we don't know the structure or family-tree equivalent when we start collecting the photographs? Suppose that the families are intermingled in ways that are not usually documented in western family trees? Suppose the kinship structure requires quite different sorts of connections to be made? All of these suppositions are probably not far from the truth. So we need a way of dynamically associating people and letting the structure emerge. We cannot be tied to a pre-determined structure.

Describing resources in the way prescribed by RDF leaves open the opportunities we need. If we say of any resource, *only something that can be said in the 'triple' format of this resource (x) has this property (y) with this value (z)*, then we can make associations as the photos are added to the collection. Instead of prescribing the associations, we are choosing to leave them open but at the same time provide enough structure for them to have the potential to be used. We provide the simple triple structure in the place of a predetermined more complex structure, and let that be defined later on.

Conclusion

At last we feel we are beginning to address the requirements at the level at which they make some sense in our context. Matchbox needs to use RDF because it is a requirement for the sort of 'doing' that a cultural heritage tool can offer its users. We cannot predict what our users will want to do, and we certainly realise now that we will not be satisfied if all we offer them is a catalogue for discovery by search, or even by browsing. We want to offer a cultural heritage tool, something that does things for them, or with which they can do things. We are aware that we will have to provide some means of accessing the power we provide, but the interface issue is not of concern here.

This paper has worked through three levels of questions that have emerged for the research team associated with the Quinkan Matchbox project. The questions parallel the freedoms offered by the advances in the relevant technologies.

Endnotes and References

1 http://www.jcu.edu.au/rockart/Matchbox

2 http://dublincore.org/

3 http://www.cimi.org/

4 http://www.w3.org/MarkUp/

5 http://www.w3.org/XML/

6 http://www.w3.org/RDF/

7 http://swordfish.rdfweb.org/discovery/2001/08/codepict/

A Prototype Digital Library For 3D Collections: Tools To Capture, Model, Analyze, and Query Complex 3D Data

Jeremy Rowe and Anshuman Razdan, Arizona State University, USA

Abstract

The Partnership for Research in Spatial Modeling (PRISM) project at Arizona State University (ASU) developed modeling and analytic tools to respond to the limitations of two-dimensional (2D) data representations perceived by affiliated discipline scientists, and to take advantage of the enhanced capabilities of 3D data that raise the level of abstraction and add semantic value to 3D data. Three-dimensional data is complex, and application of modeling and analytic techniques significantly enhances the capacity for researchers to extract meaning from 3D information. The tool prototypes simplify analysis of surface and volume using curvature and topology to help researchers understand and interact with 3D data. The tools automatically extract information about features and regions of interest to researchers, calculate quantifiable, replicable metric data, and generate metadata about the object being studied. To make this information useful to researchers, the project developed prototype interactive, sketch-based interfaces that permit researchers to remotely search, identify and interact with the detailed, highly accurate 3D models of the objects. The results support comparative analysis of contextual and spatial information, and extend research about asymmetric man-made and natural objects that can significantly extend the interactive capabilities of museums for exhibitions, education, and outreach.

Keywords: modeling, archiving, query, retrieval, three-dimensional (3D) objects

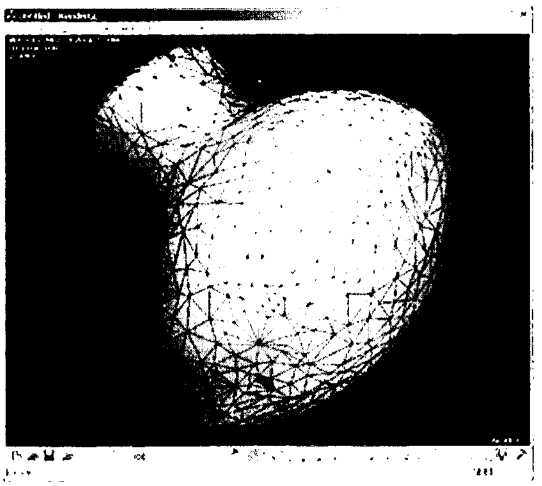

Figure 1. 3D model of Hohokam ceramic vessel.

Background

Describing, cataloguing, analyzing and organizing 3-dimensional (3D) objects have been significant and long-standing challenges to the museum community. Sketches and scientific illustrations were augmented by photography in the mid 1840s. Beginning in the 1970s, computers began to provide powerful capabilities to automate and link catalogues, to manage research data, and to combine images and text to create educational materials and programs.

Today digital museum collections and digital libraries include text, graphics, images, and increasingly, video, sound, animation, and sophisticated visual displays. Some now display three-dimensional objects and permit the user to rotate and view an image of the original object in their browser window using QuickTime, plug-ins, or custom applications. Examples range from presentation of objects for research or public access to time-lapse movies of exhibit construction and panoramas of exhibitions.

Multiple photographs and QuickTime have begun to capture representations of 3D objects, providing "rotatable" images of complex objects and environments. These photographic representations of shape can be powerful tools for interaction and education; however, the underlying images are still two-dimensional and provide insufficient information for true 3D analysis.

Though still significantly more complex and expensive than traditional photography, 3D data is becoming less costly to acquire. In addition, the number of sources of 3D data continues to increase.

Medical imaging techniques such as CAT scans and MRI yield 3D data, as can Confocal microscopes, stereophotogrammetry, satellite and remote sensors, and laser scanners. Whether extracting information from existing data or creating data for additional analysis, the availability of digital 3D representations is increasing and will continue to increase.

Once in digital form, files can be modeled and analyzed. The Partnership for Research in Spatial Modeling (PRISM) project at Arizona State University (ASU) has worked with discipline scientists in anthropology, forensics, and cellular biology to develop prototype modeling and analytic tools that enhance research by raising the level of abstraction and adding semantic value to 3D data about the natural objects being studied.

As objects become more complex in terms of variety of shape and changes in curvature, it becomes more difficult to quantify and analyze. Developing mathematical techniques to represent shape and curvature allows accurate models of the surface of 3D objects such as ceramic vessels, bones, or lithics to be created. These surface models and sophisticated mathematical tools present the ability to analyze, identify, and compare the objects that they represent. The accuracy of the measurements derived from the 3D models created equal or exceed those possible using traditional 2D tools such as calipers and rulers. In addition, measurements such as height, width, maximum height or width, surface area, or volume can be easily, consistently and accurately calculated, even for asymmetric natural objects.

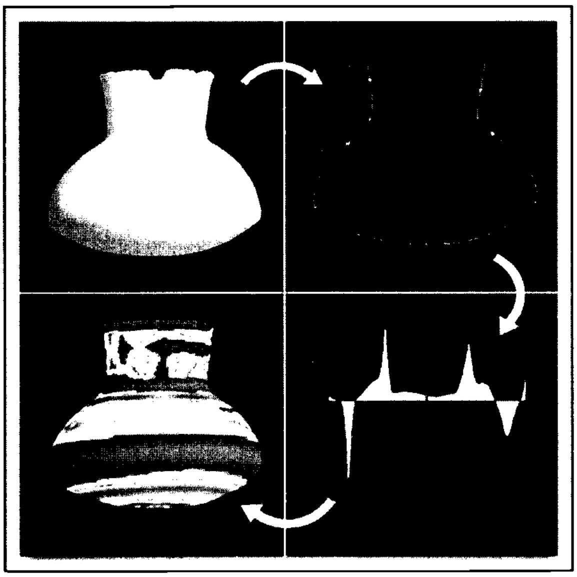

Figure 2. Examples of spatial analysis of ceramic vessel.

Use of 3D data also makes possible new measures based on topology and global or local changes in curvature that define the shape of the original object. The project built an interdisciplinary team of discipline and computer scientists and technologists to guide an interactive development processes. Research questions were initially posed by the discipline scientists; then tools and spatial modeling techniques to address them were developed by the computer scientists. With the use of mathematical models and surface and volume information, many new and powerful analytic tools become available to spatially analyze objects. For example, boundaries between surfaces can be objectively identified, small local areas of changes in curvature identified and compared, and accurate, replicable measurements calculated automatically.

Once the domain scientists link meaning to the changes in topology, shape, or curvature, a "feature" is defined. The modeling process provides an objective method to calculate physical measurements, and to consistently identify boundaries and changes that are associated with the feature, defining local areas that are of interest to researchers. Once a feature is identified, it can be described by its size, position, shape or curvature. Examples of features that can be extracted from the model data include the maximum diameter or height of a ceramic vessel.

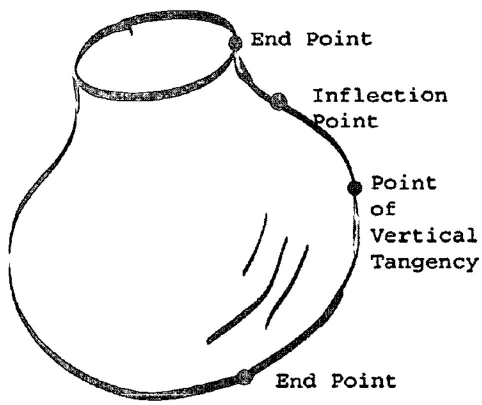

Figure 3. Components of interest for Ceramic Vessel

Features can also be mathematically abstract components of interest to the researcher, such as the base or neck of a vessel, keel of a ship, boundaries of the joint surfaces on a bone or spindles that form in the nucleus of a cell during meiosis. Often the tools developed to identify features and regions also provide additional capabilities that raise new research questions within the disciplines. Developing the tools needed to address these questions becomes a new design challenge for the computer scientists, fostering a new cycle of development. For example, ceramic analysts have found tools that identify mathematically defined features found on the vertical profile curve of a vessel; such as end points, points of vertical tangency, inflection points and corner points, as features extremely helpful in analyzing vessel shape and style. These same tools have been useful in identifying condyle surfaces of trapezia for anthropologists and forensic scientists. The tools developed to join regions of interest in the trapezia have found application to lithic tool analysis.

In addition to the tangible research benefit the tools and techniques provide, a significant result of this process has been the "cross-pollination" that has occurred as graduate students and faculty from different disciplines gravitate to a given project and explore application of tools and techniques to other discipline research.

A summary of data acquisition and analysis processes begins with laser scanning to acquire the 3D data that represents the object. Mathematical modeling is then applied to identify features and regions of interest to the domain scientists. Software tools developed by the project team generate analytic data about the original object, automatically assign metadata about spatial characteristics, and populate the database.

A visual query interface was developed to permit researchers to interact with the data using both contextual (text and numeric descriptive data) and spatial (shape and topological attribute) data. A sketch-based interface was developed to permit users to input both context and sketches to visually describe the object to initiate the search. Several text and spatial matching algorithms are used to identify and rank order objects within the database that match the search criteria.

Initial development of the digital collections focused on Classic Period (A. D. 1250 – 1450) pre-Columbian Hohokam ceramic vessels from central Arizona housed at the Archeological Research Institute at ASU. These vessels have simple, undecorated surfaces, and their analysis focuses on shape and symmetry. The level of symmetry has been a research focus as it relates to the skill of the potter and may be related to the level of time devoted to craft as a community develops and evolves over time. Since even the best hand-made pots are asymmetric, the traditional measure of symmetry, the profile curve, can vary dramatically depending on the orientation of the vessel. Multiple photographs used to create a QuickTime view of the vessel would assist researchers in visual analysis, but not in more detailed measurement. By scanning and creating a 3D model of the vessel, researchers can perform detailed, objective analysis of the shape and symmetry using tools to compare local and overall curvature, inflection points (changes in curvature from convex to concave), corner points, and calculated measures such as surface area.

Methods

Metadata Schema and Organizational Structure

One of the greatest challenges in an interdisciplinary research effort is coordinating expectations among team members, and developing communication processes that bridge conceptual, strategic, and linguistic differences across the disciplines. An iterative process was developed to share research questions, tools and intellectual approaches across disciplines at project meetings. The results were a gradual bonding of researchers, development of a shared vocabulary, and substantial interaction about potential research issues and approaches. These efforts provided a foundation for the initial modeling and analysis, and for developing the metadata structure needed to organize data for storage, analysis, and query.

A conceptual goal of the metadata component of the project was to develop an extensible schema structure that could accommodate adding new types of objects as the project continued to evolve. An object class was defined as the master class docu-

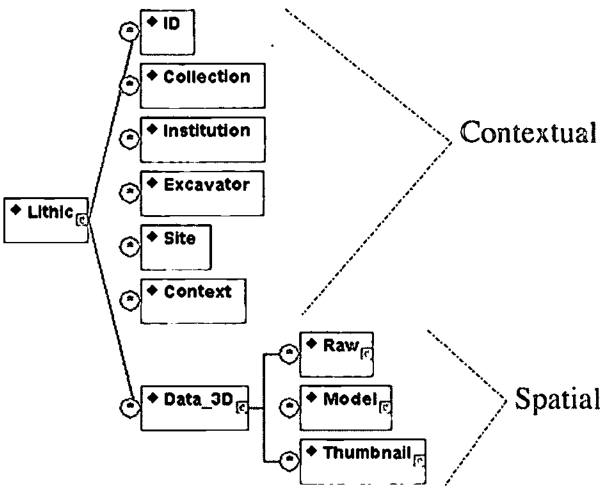

Figure 4. Example of metadata structure.

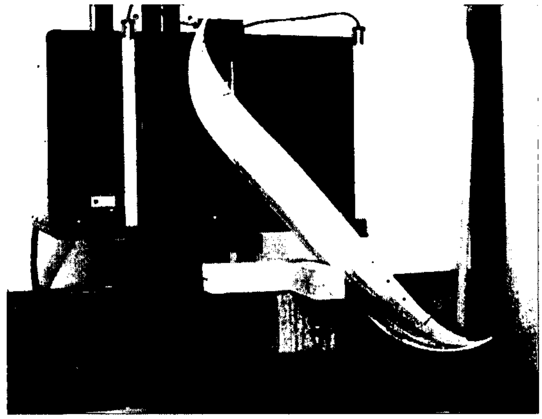

Figure 5. Laser scanning ship model

ment type definition (DTD) for each item in the digital library database. For the 3DK digital library project, all of the additional descriptive data about each object was defined and organized as contextual or spatial classes.

Contextual types define text and metric information about the object. This context class includes subclasses for metadata associated with objects as they are acquired, processed, and archived; such as type, item name, catalogue number, collection, provenance, etc. At this phase of the project, these fields were primarily determined by existing descriptive data elements, though efforts were made to design a schema structure that would accommodate adding new object types as necessary. To date, several iterations to refine the schema model to function effectively across object types have been completed.

Spatial data types define the 3D attributes of the object, including raw data, thumbnails, models, and calculated or derived data about the topology, shape, and composition of the object. Use of common descriptive components and geometric elements as new object types are added will permit shared use of the modeling and analysis tools across classes of objects. The project goal is to develop standards for description and organization that permit automated cataloguing and population of data as objects are scanned and processed for entry into the database.

Due to familiarity and availability of resources, an SQL database was used to store the contextual and spatial data. Fields were assigned to each data element, and large spatial data files were stored as hyperlinks. Generally accepted data formats such as binary, PLY, HTML, and XML have been used to make data accessible and simplify migration and access to the data over time.

Scanning to Acquire 3D Data

The PRISM Digital Library project uses two Cyberware scanners, the M15 and 3030, to scan and capture 3D data describing ceramic vessels, bones, and other objects up to roughly a 30î maximum dimension. Each object is scanned by a laser which captures spatial data (x, y, z) values for each point. The scanners capture line-of-site data, so each object must be scanned, then rotated, and scanned again to capture additional data. This process is repeated until sufficient scans are obtained to combine to create a point cloud model to document the surface.

The Model 15 laser digitizer captures surface data points less than 300 microns (0.3mm) apart, producing high-density triangular meshes with an average resolution of over 1000 points per cm2. The digitized data generated by the scanner is composed of thousands of (x, y, z) coordinates that describe a point cloud that represents the surface of the object scanned. Further analysis requires generating a surface model from the point cloud.

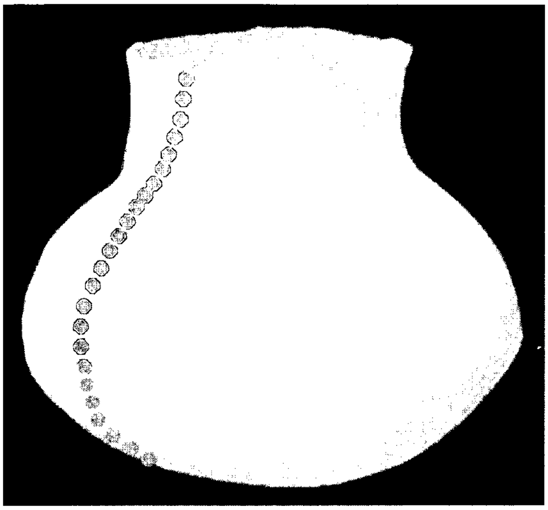

Figure 6. Representation of scanning of ceramic vessel

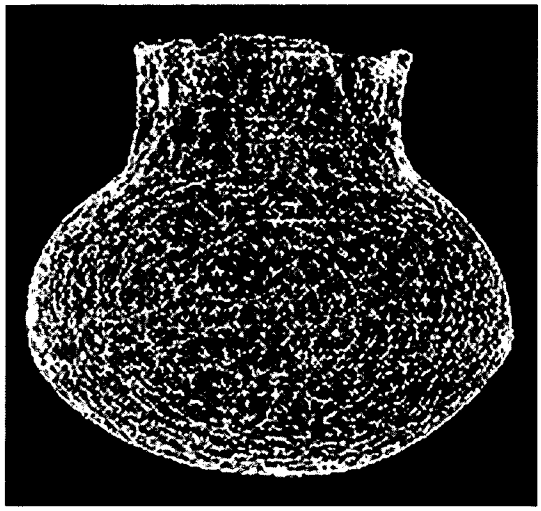

Figure 7. Point cloud of ceramic vessel combined from multiple scans.

Modeling techniques are used to create an actual measurable surface that represents the original object. In addition to the triangle meshes, PRISM software can represent these surfaces as Non-Uniform Rational B-spline (NURB) or subdivision surfaces (Bernadini et al., 98; Razdan et al., 98; Farin, 01, Farin, 02). NURB representation provides the capability to assess curvature distribution in complex objects; such as identification of the joint surfaces from scanned data of a bone.

The accurate model of the object that results from this process provides the data and conceptual framework needed for objective, replicable analysis of surface and volume attributes of the objects under study.

Extracting Features and Identifying Regions of Interest

Once the geometric structure has been obtained, the next step is to identify features and regions of interest to the discipline researchers. Ceramicists look for shape, symmetry, and curvature, cellular biologists look for structure of bio-molecular machines inside a cell, forensic anthropologists look at shape, and surface comparisons. A number of 3D modeling and analytic algorithms have been combined, and new techniques developed to segment the geometric structure into regions, and to identify meaningful features

The nontrivial challenge has been to translate the features of interest to the discipline scientists into mathematically definable terms. For example, the transition between the neck and body of a vessel can be described mathematically as an inflection point and the maximum width of a vessel by its greatest diameter. Crosswalks of definitions to help translate terms and permit mapping mathematical concepts on to features meaningful to the discipline scientists have been developed by the project team. The 3D data permits accurate maximum and minimum measurements to be identified, as well as allowing calculation of complex metric and descriptive data that are extremely difficult to obtain using 2D representations, linear measurements, and traditional measuring tools, particularly for naturally asymmetric or man-made objects such as ceramics.

The second program developed is Region Editor that calculates more complex information about the object and its component features such as total object volume, absolute object symmetry, the area of surfaces identified, and the average angle at which surfaces intersect. Several of these measures are extremely difficult to determine accurately using traditional techniques, such as tape and caliper, particularly for asymmetrical objects. The Region Editor also permits researchers to add contextual information such as technical data about the scan, image processing that has been used, provenance,

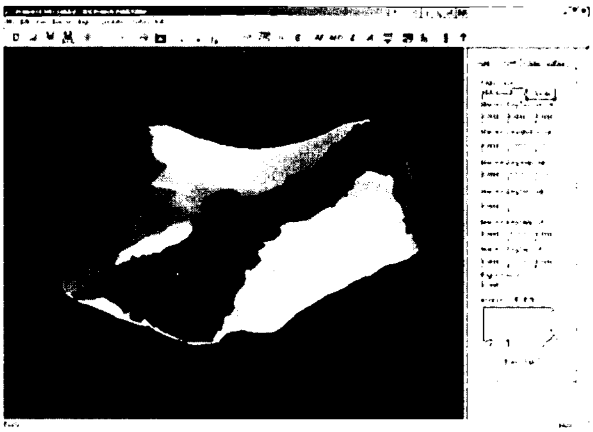

Figure 8. Region editor applied to trapezium data model.

or collection to the 3D data. The final action of the Region Editor is to create the metadata or XML file associated with the 3D data for archiving.

Interacting with the Data: the User Interface

A primary design problem was how to accept input to support searches for both contextual and spatial variables. An interdisciplinary "visual query interface" team guided research into interface design, identified desired capabilities, developed the interface, and coordinates ongoing revision based on evaluation data.

The PRISM team chose to design separate contextual and spatial input areas in the interface screen. Textual data is input or selected from pull-down menus to query existing descriptive catalogues or databases. Search criteria can include metadata such as name, type or number of the item, collection, or other catalogue information about the object. This input area also permits the user to limit search by provenance by limiting the search to a specific col-

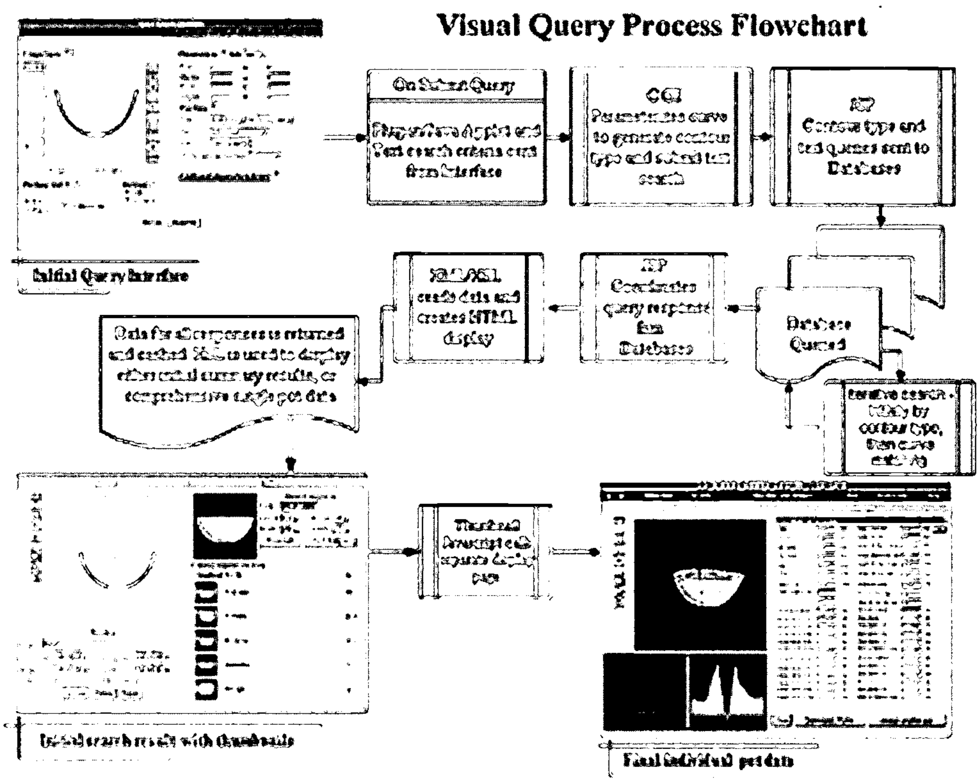

Figure 9. Diagram of Visual Query Interface and Search Process.

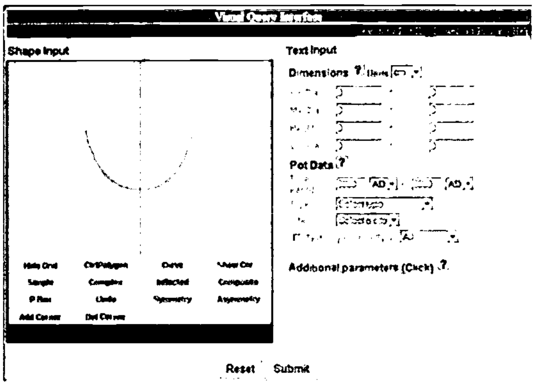

Figure 10. Prototype profile-based visual query interface for searching ceramic vessels

lection, or by measurements such as height, width, or maximum or minimum diameter.

The most interesting interface design challenge was to accommodate the input to query spatial data and to identify matching 3D shapes (Sakurai and Gossard, 88; Osada et al 2001, 2002; Razdan et al. 2001).To simplify the initial development of the prototype, and to mirror the 2D shape profiles of the ceramic vessels familiar to anthropologists, the interface model uses an interactive vessel profile graphic representation to define the spatial search component. A gridded area is used to present a sample of a profile curve selected from the menu, or permit the researcher to draw a profile to be searched. Using the mouse and tool palette, the user can interactively create or manipulate the shape until it represents the desired vessel. Initially developed as a Netscape plug-in, the sketch interface has been converted into a Java applet to support multiple browsers and platforms.

After descriptive information about context and shape has been entered, the query is submitted.The descriptive and spatial information are separated, and the multiple database queries are coordinated by project software.The contextual component of the query is handled as a conventional text and numeric database search.The spatial search uses a variety of size, shape, and curve matching algorithms developed by the project team to identify and locate similarities within the databases.

During search and analysis of potential matches, intelligent filtering techniques are used to limit the search pool. Initially simple text, metric, or gross spatial classification criteria are used to identify possible matches from the database and reduce the search domain The search progressively applies increasingly more complex algorithms to the shrinking pool of potential matches. This process minimizes computational load and search time while accurately identifying all objects that match the search criteria.

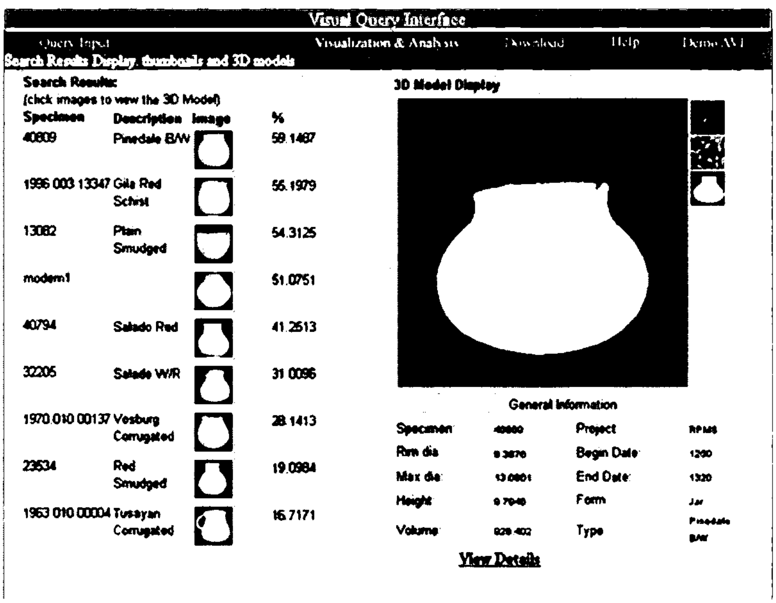

Figure 11. Ranked search results.

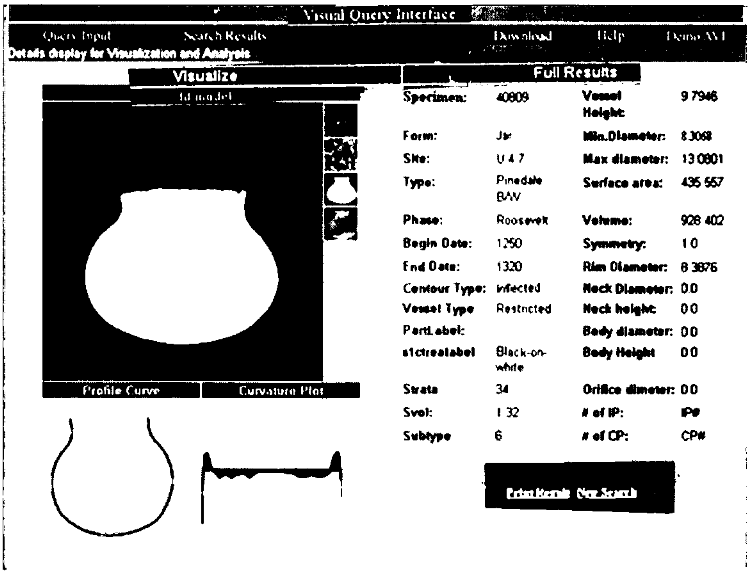

Figure 12. Detailed search results

Another algorithm ranks the query results by descriptive and spatial similarity to the query image. Query response information is presented sequentially over several screens, each providing an additional level of information about the selected objects.The first screen displays thumbnail images and brief descriptions of the top search results. Also presented is a large 3D display of the top search result, along with more detailed descriptive and calculated information.The 3D model can be displayed as a point cloud, wire frame, or full shaded surface representation at the discretion of the researcher. Using the mouse, the model can be rotated and viewed from any angle.Selecting a thumbnail of another search result from the queue of search results will replace its model in the 3D display window.

If more detailed descriptive information is desired, a third window is available to display the 3D model and two additional analytic tools - a profile curve and curvature plot, and additional descriptive data about the object.A fourth window can be selected to provide access to the complete descriptive and calculated data available.

Significant effort has been given to adapting the interface design to accommodate differences in contextual data and analytic tools across different classes of objects.The object type metadata can be used to select the customized search template with fields for the contextual and spatial data appropriate for the object.The visual query interface team developed training materials to guide new users and evaluation instruments to obtain formative guidance from users.

Evaluation

Several techniques were used to evaluate and guide the development of the project. In addition to general meetings and team building activities, process mapping and interviews with project team members provided qualitative and quantitative input to help build communication among researchers in the team.

Initial evaluation input regarding interface components and design were obtained from the roughly 25 project team members.The current version of the interface was used and assessed by the entire group at general and visual query interface team meetings throughout its development.The designs were critiqued, limitations identified, additional desired capabilities described, development challenges identified, and component work delegated to project teams.

Several evaluation sessions were held to obtain input from faculty and student researchers outside of the team.After initial orientation, research problems were posed to the evaluation groups, and us-

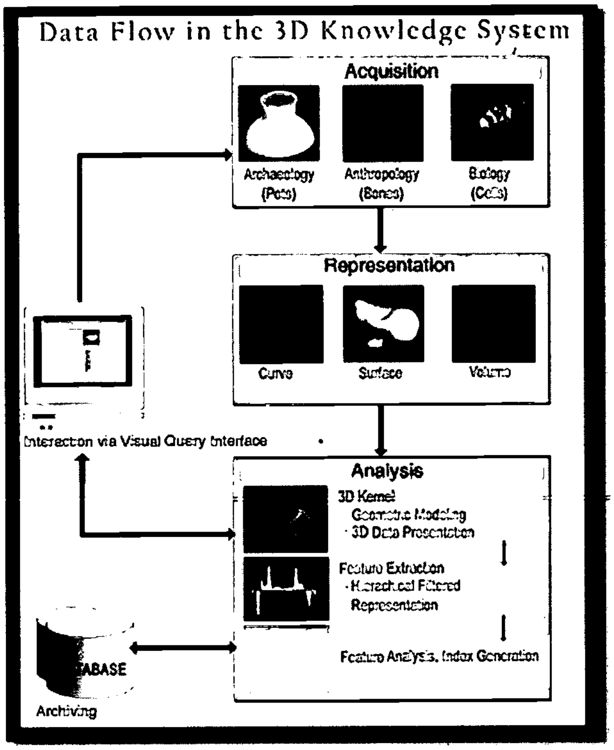

Figure 13. Overview of prototype digital library process.

ers used the interface to locate individual target objects by context, shape or size. Users were encouraged to explore the 150ceramic vessels in the test database and comment on the clarity, scope, and ease of use of the interface. A revision cycle followed each evaluation.

Discussion

One of the pleasant surprises during this project has been the ease of extending the modeling and analytic tools developed for one specific discipline to other research domains, as well as the growth of the tools for surface and volume modeling and analysis. The improvements that have resulted from the iterative process of identifying a domain research question, developing an application tool, deployment, analysis of potential applications across other research domains, and identification of new research questions have generated significant progress in developing modeling and analytic tools applicable to 3D data.

As 3D data acquisition tools become more affordable and readily available, the amount of 3D data that must be described, stored and displayed will grow dramatically. Accommodating this huge data management challenge will require development of standards and tools to begin to analyze and add meaning to the data.

The spatial and volume modeling and analytic tools developed by the project team permit discipline researchers to quantify and accurately replicate measurements of complex 3D objects. The feature and region recognition capabilities assist in visualizing complex, abstract concepts of interests to discipline researchers.

The ability to generate and analyze accurate representations of 3D objects has many potential uses in the research and museum communities. Scanned and modeled artifacts can permit second-best access to the original objects where access is restricted, such as by country of origin, or by delicacy or condition of the object.

Since the models are scale-independent and can vary in size or magnification, they permit detailed analysis augmenting visual tools such as dissecting scopes or microscopes for small objects, and providing new perspectives for viewing and analyzing large scale objects.

In addition to analysis of individual objects, the tools can be used to compare objects, and one application currently under development is lithic refitting, though the time and effort currently involved limit the practicality of this approach.

These modeling and analytic techniques appear to have many other applications for museums and collection management. The representations are accurate enough to support condition analysis to determine changes over time by comparing sequential scans. For Native American materials, another potential application is documenting repatriated objects to identify artifacts that reappear on the market at some point in the future.

Several efforts are underway or are planned by the PRISM team to further extend the capabilities of the tools developed and their application to domain research. In terms of infrastructure, the move from custom plug-ins to Java will simplify deployment.

We are exploring alternatives to the currently used SQL database, such as object-oriented databases. Another effort to improve searching is a pilot XML search protocol developed by the National Science Foundation Biological Databases and Informatics project at Arizona State University (BDI) research project, in conjunction with the ASU Long Term Ecological Research (LTER) Metadata Committee and the Knowledge Network for Biocomplexity (KNB) Project at the National Center for Environmental Analysis and Synthesis (NCEAS). The "Xanthoria" metadata query system developed by this project team uses SOAP (Simple Object Access Protocol) to send XML query requests and responses, and supports simultaneous Web-based querying of distributed, structurally different metadata repositories.

The spatial analytic tools continue to develop as improvements are made in the feature extraction and region editing applications and more powerful techniques are developed to compare curvature, identify matches and rank search results. Key to these efforts are the expanding partnerships with other research areas with their own unique modeling and visualization needs. Included to date are more complex anatomical data from CAT scanners and MRI, cloud formation pattern recognition, geological erosion, and identification of targets within complex, noisy environmental data.

Interface design continues to evolve. The project is evaluating models developed for 3D query and display by other projects including:

- Princeton 3D Models Search Engine using Takeo Igarashi's Teddy 3D sketch interface

- National Center for Biotechnology Information (NCBI) Cn3D Genetic viewer

The goal is to develop realistic 3D interface models that permit the researcher to sculpt the query in 3D space. Additional analytic tools are also being developed, such as planar overlays to visualize and objectively compare joint surfaces of bones. Techniques to bookmark searches to permit replication and simplify comparison of objects within the databases are a significant challenge for dynamic data sets in a client-server environment and are also being explored. A complex variation of bookmarks involves providing a replicable trail for researchers using the region editor and additional analytic tools such as the planar overlay to interact with the data and create their own interpretive models. Creating storage techniques for these derived, researcher defined or modeled data, and managing "version control" to permit replication and deconstruction of the analysis, is another challenge.

User evaluation of the current interface layout, color palette and design continues, using both surface and volume model data. In addition to initially developing specific bone or ceramic vessel interfaces for the different research domains, the project is working to identify commonalities and conventions to develop a unified interface model. This common design appears to be possible in initial query interface screens, with differentiation of interface display occurring as objects are identified, search results are returned, and researchers drill down into object data that may vary across disciplines.

Acknowledgements

This work was supported in part by the National Science Foundation (grant IIS-9980166) and funding from the Vice Provost for Research and Economic Development at Arizona State University. The authors would like to thank all of the collaborators that make up the Partnership for Research in Spatial Modeling (PRISM) team, particularly Anshuman Razdan, Gerald Farin, Daniel Collins, Peter McCartney, Matthew Tocheri, Mary Zhu, Mark Henderson, Arleyn Simon, Mary Marzke David Capco. For more information on the 3D Knowledge project visit http://3dk.asu.edu.

References

Amresh, A., G. Farin, & A. Razdan. Adaptive Subdivision Schemes for Triangular Meshes, will appear in "Hierarchical and Geometric Methods in Scientific Visualization", edited by G. Farin, H. Hagen, B. Hamann, Springer-Verlag, 2002 (in print).

Bernadini, F., J. Mittleman, H. Rushmeier, C. Silva, G. Taubin. The Ball-Pivoting Algorithm for Surface Reconstruction, IEEE Transactionson Visualization and Computer Graphics, Vol. 5, No. 4, October/December 1999.

Council for Preservation of Archeological Records (COPAR) http://copar.asu.edu/

Dublin Core Metadata Initiative http://dublincore.org/documents/2000/07/11/dcmes-qualifiers/

Farin , G. "Curves and Surfaces for CAGD", 5th ed., Morgan-Kaufmann, 2001.

Farin, G. History of Curves and Surfaces in CAGD. In: Handbook of CAGD, G. Farin, M.S. Kim, J. Hoschek (eds), Elsevier, 2002.

Hamann, B,, Jean, A. Razdan. CAGD Techniques in the Control of Surface Grid Generation. In: Thompson, J.F., N.P Weatherill, & B.K. Soni, (eds.), Handbook of Grid Generation, CRC Press, Inc., Boca Raton, Fla. pp 29.1-26, 1997

Knowledge Network for Biocomplexity http://knb.ecoinformatics.org/

Long Term Ecological Research project at Arizona State University http://caplter.asu.edu/

Mangan, A. and R. Whitaker. Partitioning 3D Surface Meshes Using Watershed Segmentation. IEEE Transactions on Visualization and Computer Graphics. Vol.5, No. 4, Oct-Dec 1999.

National Center for Biotechnology Information (NCBI) Cn3D Genetic viewer http://www.ncbi.nlm.nih.gov/Structure/CN3D/cn3d.shtml

National Center for Environmental Analysis and Synthesis http://cochise.asu.edu/bdi/Subjects/Xanthoria/index.htm

Osada, R., T. Funkhouser, B. Chazelle, and David Dobkin, Shape Distributions, to appear in ACM Transactions on Graphics, 2001.

Osada, R., T. Funkhouser, B. Chazelle, and D. Dobkin, Matching 3D Models with Shape Distributions, Shape Modeling International, Genova, Italy, May 2001.

Princeton 3D Models Search Engine using Takeo Igarashi's Teddy 3D sketch interface – http://www.cs.princeton.edu/gfx/proj/shape

Partnership for Research in Spatial Modeling at Arizona State University http://3DK.ASU.EDU

A. Razdan and Myung Soo Bae, A Hybrid Approach to Feature Segmentation, in preparation.

Razdan, A., B. Steinberg, & G. Farin. From Digitized Data to NURB Surface Meshes: Proceedings of the International Conference of Rapid Prototyping and Manufacturing, pp 749-754, Beijing, China, 1998.

Razdan, A., D. Liu, M. Bae, M. Zhu, G. Farin, A. Simon, M. Henderson. Using Geometric Modeling for Archiving and Searching 3D Archaeological Vessels. CISST 2001 June 25-28, 2001, Las Vegas.

Rowe, J. Developing a 3D Digital Library for Spatial Data: Issues Identified and Description of Prototype, RLG DigiNews October 2002, http://www.rlg.org/preserv/diginews/diginews6-5.html#feature1

Sakurai, H., and D. Gossard, Shape Feature Recognition from 3D Solid Models. In ASME Computers in Engineering, San Francisco, 1988.

Schurmans, U., A. Razdan, A. Simon, P. McCartney, M. Marzke, D. Van Alfen, G. Jones, J. Rowe, G. Farin, D. Collins, M. Zhu, D. Liu, and M. Bae, "Advances in Geometric Modeling and Feature "Extraction on Pots, Rocks and Bones for Representation and Query via the Internet," proceedings Computer Applications in Archaeology (CAA), 2001.

Simon, A., D. Van Alfen, A. Razdan, G. Farin, M. Bae, and J. Rowe, "3D Modeling for Analysis and Archiving of Ceramic Vessel Morphology: A Case Study from the American Southwest," Proceedings of the 33rd International Symposium on Archaeometry. Geoarchaeological and Bioarchaeological Studies, Vrije Universiteit, Amsterdam, 2002.

Interfaces

The More You Look the More You Get: Intention-based Interface using Gaze-tracking

Slavko Milekic, The University of the Arts, USA

http://www.uarts.edu

Abstract

Only a decade ago eye- and gaze-tracking technologies using cumbersome and expensive equipment were confined to university research labs. However, rapid technological advancements (increased processor speed, advanced digital video processing) and mass production have both lowered the cost and dramatically increased the efficacy of eye- and gaze-tracking equipment. This opens up a whole new area of interaction mechanisms with museum content. In this paper I will describe a conceptual framework for an interface, designed for use in museums and galleries, which is based on non-invasive tracking of a viewer's gaze direction. Following the simple premise that prolonged visual fixation is an indication of a viewer's interest, I dubbed this approach intention-based interface.

Keywords: eye tracking, gaze tracking, intention-based interface

Figure 1. Comparison of human and non-human eye (chimpanzee). Although many animals have pigmentation that accentuates the eyes, the visible white area of human eye makes it easier to interpret the gaze direction

Introduction

In humans, gaze direction is probably the oldest and earliest *means of communication at a distance*. Parents of young infants are often trying to 'decode' from an infant's gaze direction the needs and interest of their child. Thus, gaze direction can be viewed as a first instance of *pointing*. A number of developmental studies (Scaife and Bruner 1975; Corkum and Moore, 1988; Moore 1999) show that even very young infants actively follow and respond to the gaze direction of their caregivers. The biological significance of eye movements and gaze direction in humans is illustrated by the fact that humans, unlike other primates, have visible white area (sclera) around the pigmented part of the eye (iris, covered by transparent cornea, see Figure 1). This makes even discrete shifts of gaze direction very noticeable (as is painfully obvious in cases of 'lazy eye').

Eye contact is one of the first behaviors to develop in young infants. Within the first few days of life, infants are capable of focusing on their caregiver's eyes (Infants are physiologically shortsighted with the ideal focusing distance of 25-40 cm. This distance corresponds to the distance between the mother's and infant's eyes when the baby is held at the breast level. Everything else is conveniently a blur)[1]. Within the first few weeks, establishing eye contact with the caregiver produces a smiling reaction (Stewart & Logan, 1998). Eye contact and gaze direction continue to play a significant role in social communication throughout life. Examples include:

- regulating conversation flow;
- regulating intimacy levels;
- indicating interest or disinterest;
- seeking feedback;
- expressing emotions;
- influencing;
- signaling and regulating social hierarchy;
- indicating submissiveness or dominance;

Thus, it is safe to assume that humans have a large number of behaviors associated with eye movements and gaze direction. Some of these are innate (orientation reflex, social regulation), and some are learned (extracting information from printed text, interpreting traffic signs).

Our relationship with works of art is essentially a social and intimate one. In the context of designing a gaze tracking-based interface with cultural heritage information, innate visual behaviors may play a significant role precisely because they are social and emotional in nature and have the potential to *elicit a reaction* external to the viewer. In this paper I will provide a conceptual framework for the design of gaze-based interactions with cultural heritage information using the digital medium. Before we proceed, it is necessary to clarify some of the basic physiological and technological terms related to eye- and gaze-tracking.

Eye Movements and Visual Perception

While we are observing the world, our subjective experience is that of a smooth, uninterrupted flow of information and a sense of the wholeness of the visual field. This, however, contrasts sharply with what actually happens during visual perception. Our eyes are stable only for brief periods of time (200-300 milliseconds) called *fixations*. Fixations are interspersed by rapid, jerky movements called *saccades*. During these movements no new visual information is acquired. Furthermore, the information gained during the periods of fixations is clear and detailed only in a small area of the visual field – about 2 of visual angle. Practically, this corresponds to the area covered by one's thumb at arm's length. The rest of the visual field is fuzzy but provides enough information for the brain to plan the location of the next fixation point. The problems that arise because of the discrepancy between our subjective experience and the data gained by using eye-tracking techniques can be illustrated by the following example:

(221) (268) (292) (197) (204) (177) (156)

The horse raced past the barn fell.

1 6 2 3 7 4 5

The sentence above is a classical example of a "garden path" sentence that (as you probably have experienced) initially leads the reader to a wrong interpretation (Bever, 1970). The eye-tracking data provide information about the sequence of fixations (numbered 1 to 7) and their duration in milliseconds. The data above provide some clues about the relationship between visual analysis during reading and eye movements. For example, notice the presence of two *retrograde saccades* (numbered 6 and 7) that happened after initial reading of the sentence. They more than double the *total fixation time* of the part of the sentence necessary for disambiguation of its meaning. Nowadays there is a general consensus in the eye-tracking community that the number and the duration of fixations are related to the cognitive load imposed during visual analysis.

Path (1) corresponds to free exploration. Path (2) was obtained when subjects were asked to judge the material status of the family, and path (3) when they were asked to guess the age of different individuals. Partially reproduced from Yarbus, A. L. (1967)

Eye-tracking studies of reading are very complex but have the advantage of allowing fine control of different aspects of the visual stimuli (complexity, length, exposure time, etc.). Interpretation of eye movement data during scene analysis is more com-

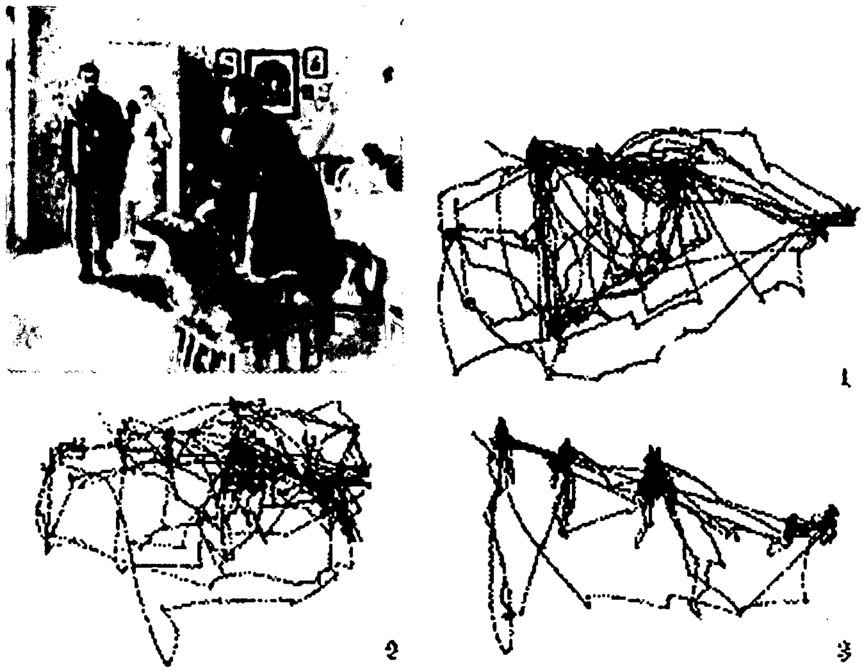

Figure 2. Illustration of differences in gaze paths while interpreting I. Repin's painting "They did not expect him."

plicated because visual exploration strategy is heavily dependent on the *context of exploration*. Data (Figure 2) from an often-cited study by Yarbus (1967) illustrate differences in visual exploration paths during interpretation of Ilya Repin's painting "They did not expect him, or "the unexpected guest"

Brief History of Eye- and Gaze-Tracking

The history of documented eye- and gaze-tracking studies is over a hundred years old (Javal, 1878). It is a history of technological and theoretical advances where progress in either area would influence the other, often producing a burst of research activity that would subsequently subside due to the uncovering of a host of new problems associated with the practical uses of eye-tracking.

Not surprisingly, the first eye-tracking studies used other humans as tracking instruments by utilizing strategically positioned mirrors to infer gaze direction. Experienced psychotherapists (and socially adept individuals) still use this technique, which, however imperfect it may seem, may yield a surprising amount of useful information. Advancements in photography led to the development of a technique based on capturing the light reflected from the cornea on photographic plate (Dodge & Cline, 1901). Some of these techniques were fairly invasive, requiring placement of a reflective white dot directly onto the eye of the viewer (Jud, McAllister & Steel, 1905) or a tiny mirror, attached to the eye with a small suction cup (Yarbus, 1967). In the field of medicine a technique was developed (electro-oculography, still in use for certain diagnostic procedures) that allowed registering of eyeball movements using a number of electrodes positioned around the eye. Most of the described techniques required the viewer's head to be motionless during eye tracking and used a variety of devices like chin rests, head straps and bite-bars to constrain the head movements. The major innovation in eye tracking was the invention of a head-mounted eye tracker (Hartridge & Thompson, 1948). With technological advances that reduced the weight and size of an eye tracker to that of a laptop computer, this technique is still widely used.

Most eye tracking techniques developed before the 1970s were further constrained by the fact that data analysis was possible only *after* the act of viewing. It was the advent of mini- and microcomputers that made possible real-time eye tracking. Although widely used in studies of perceptual and cognitive processes, it was only with the proliferation of personal computers in the 1980s that eye tracking was applied as an instrument for the evaluation of human-computer interaction (Card, 1984). Around the same time, the first proposals for the use of eye tracking as a means for user-computer communication appeared, focusing mostly on users with special needs (Hutchinson, 1989; Levine, 1981). Promoted by rapid technological advancements, this trend continued, and in the past decade a substantial amount of effort and money was devoted to the development of eye- and gaze-tracking mechanisms for human-computer interaction (Vertegaal, 1999;Jacob, 1991; Zhai, Morimoto & Ihde, 1999). Detailed analysis of these studies is beyond the scope of this paper, and I will refer to them only insofar as they provide reference points to my proposed design. Interested readers are encouraged to consult several excellent publications that deal with the topic in much greater detail (Duchowsky, 2002; Jacob, Karn, 2003 /in press/).

Eye and Gaze Tracking in a Museum Context

The use of eye and gaze tracking in a museum context extends beyond interactions with the digital medium. Eye tracking data can prove to be extremely useful in revealing how humans observe real artifacts in a museum setting. The sample data and the methodology from a recent experiment conducted in the National Gallery in London (in conjunction with the Institute for Behavioural Studies) can be seen on the Web. Although some of my proposed gaze-based interaction solutions can be applied to the viewing of real artifacts (for example, to get more information about particular detail that a viewer is interested in), the main focus of my discussion will be on the development of affordable and intuitive gaze-based interaction mechanisms with(in) the digital medium. The main reason for this decision is the issue of accessibility to cultural heritage information. Although an impressive 4000 people participated in the National Gallery experiment, they *all had to be there* at certain time. I am not disputing the value of experiencing the real artifact, but the introduction of the digital medium has dramatically shifted the role of museums from *collection & preservation* to *dissemination & exploration*. Recent advancements in Web-based technologies make it possible for museums to develop tools (and social contexts) that allow them to serve as centers of knowledge transfer for both local and virtual communities. My proposal will focus on three issues:

1. problems associated with use of gaze tracking data as interaction mechanism;
2. conceptual framework for the development of gaze-based interface;
3. currently existing (and affordable) technologies that could support non-intrusive eye and gaze tracking in a museum context.

1. Problems associated with gaze tracking input as an interaction mechanism

The main problem associated with use of eye movements and gaze direction as an interaction mechanism is known in the literature as "Midas touch" or "the clutch" problem (Jacob, 1993). In simple terms, the problem is that if looking at something should trigger an action, one would be triggering this action even by just observing a particular element on the display (or projection). The problem has been addressed numerous times in literature, and there are many proposed technical solutions. Detailed analysis and overview of these solutions is beyond the scope of this paper. I will present here only a few illustrative examples.

One of the solutions to the *Midas Touch* problem, one developed by Risø National Research Laboratory, was to separate the gaze-responsive area from the observed object. The switch (aptly named EyeCon) is a square button placed next to the object that one wants to interact with. When the button is focused (ordinarily for half a second), it 'acknowledges' the viewer's intent to interact with an animated sequence depicting a gradually closing eye. The completely closed eye is equivalent to the pressing of a button (see Figure 3).

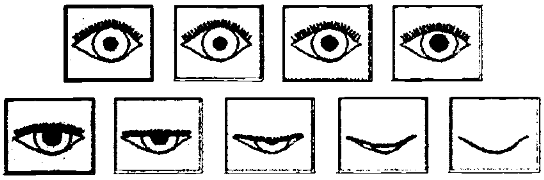

Figure 3. An EyeCon activation sequence. Separating the control mechanism from interactive objects allows natural observation of the object (image reproduced from Glenstrup, A.J., Engell-Nielsen, T., 1995)

One of the problems with this technique comes from the very solution — it is the *separation of selection and action.* The other problem is the *interruption of the flow of interaction* – in order to select (interact with) an object, the user has to focus on the action button for a period of time. This undermines the unique quality of gaze direction as the fastest and natural way of pointing and selection (focus).

Another solution to the same problem (with very promising results) was to provide the 'clutch' for interaction through another modality - voice (Glenn, Iavecchia, Ross, Stokes, Weiland, Weiss, Zakland 1986) or manual (Zhai, Morimoto, Ihde 1999) input.

The second major problem with eye movement input is the sheer volume of data collected during eye-tracking and its meaningful analysis. Since individual fixations carry very little meaning on their own, a wide range of eye tracking metrics has been developed in the past 50 years. An excellent and very detailed overview of these metrics can be found in Jacob (2003/in print). Here, I will mention only a few that may be used to infer viewer's interest or intent:

- *number of fixations*: a concentration of a large number of fixations in a certain area may be related to a user's interest in the object or detail presented in that area when viewing a scene (or a painting). Repeated, retrograde fixations on a certain word while reading text are taken to be indicators of increased processing load (Just, Carpenter 1976).

- *gaze duration*: gaze is defined as a number of *consecutive* fixations in an area of interest. Gaze duration is the total of fixation durations in a particular area.

- *number of gazes*: this is probably a more meaningful metric than the number of fixations. Combined with gaze duration, it may be indicative of a viewer's interest.

- *scan path*: the scan path is a line connecting consecutive fixations (see Figure 2, for example). It can be revealing of a viewer's visual exploration strategies and is often very different in experts and novices.

The problem of finding the right metric for interpretation of eye movements in a gallery/museum setting is more difficult than in a conventional research setting because of the complexity of the visual stimuli and the wide individual differences of users. However, the problem may be made easier to solve by dramatically constraining the number of interactions offered by a particular application and making them correspond to the user's expectations. For example, one of the applications of the interface I will propose is a simple gaze-based browsing mechanism that allows the viewer to quickly and effortlessly leaf through a museum collection (even if he/she is a quadriplegic and has retained only the ability to move the eyes).

II. Gaze-based interface for museum content

Needless to say, even a gaze-based interface that is specifically designed for museum use has to provide a solution for general problems associated with the use of eye movement-based interactions. I will approach this issue by analyzing three different strategies that may lead to the solution of the *Midas touch* problem. These strategies differ in terms of the of the interaction mechanism, as it relates to:

- time

- location, and

- user action

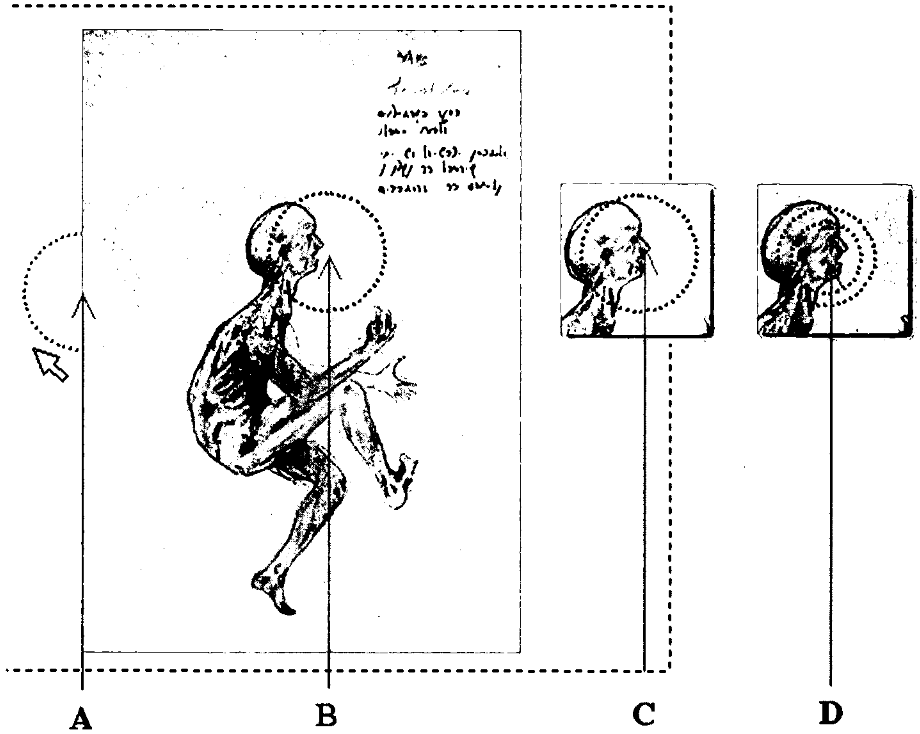

Figure 4. The cursor changes at position (A) into focus area indicating that the object is 'hot'.

It is clear that *any* interaction involves time, space and actions, so the above classification should be taken to refer to the key component of the interface solution. Each of these solutions has to accommodate two modes of operation:

- the observation mode, and
- the action (command) mode

The viewer should have a clear indication as to which mode is currently active, and the interaction mechanism should provide a way to switch between the modes quickly and effortlessly.

Time-based interfaces

At first glance, a time-based interface seems like a good choice (evident even for myself when choosing the title of this paper). An ideal setup (for which I will provide more details in the following sections) for this type of interface would be a high-resolution *projection* of a painting on the screen with an eye-tracking system concealed in a small barrier in front of the user. An illustration of a time-based interaction mechanism is provided in Figure 4. The gaze location is indicated by a traditional cursor as long as it remains in a non-active (in this case, outside of the painting) area. When the user shifts the gaze to the gaze-sensitive object (painting), the cursor changes its shape to a faint circle, indicating that the observed object is aware of the user's *attention*. I have chosen the circle shape because it does not interfere with the viewer's observation, even though it clearly indicates potential interaction. As long as the viewer continues visual exploration of the painting there is no change in status. However, if the viewer decides to focus on a certain area for a predetermined period of time (600 ms), the cursor/ circle starts to shrink (zoom), indicating the beginning of the *focusing* procedure.

Position (B) marks the period of relative immobility of the cursor and the beginning of the focusing procedure. Relative change in the size of the focus area (C) indicates that focusing is taking place. The appearance of concentric circles at time (D) indicates imminent action. The viewer can exit the focusing sequence at any time by moving the point of observation outside of the current focus area.

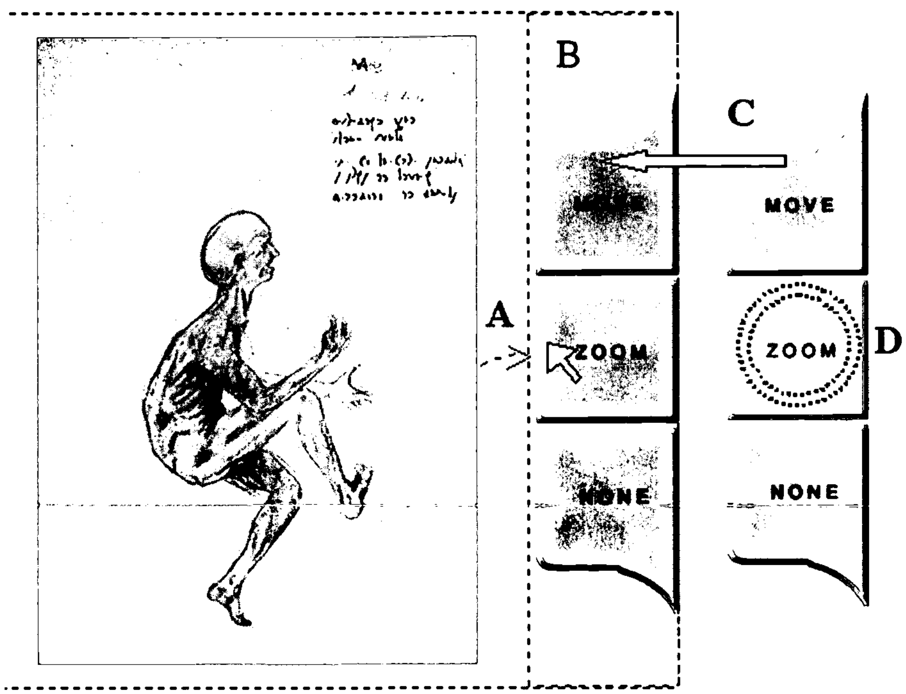

Figure 5. Movement of the cursor (A) into the gaze-sensitive area (B) slides the action palette (C) into view .

If the viewer continues to fixate on the area of interest, the focusing procedure continues for the next 400 milliseconds, ending with a 200 millisecond long signal of *imminent action*. At any time during the focusing sequence (including the imminent action signal), the viewer can return to observation mode by moving the gaze away from the current fixation point. In the scenario depicted above (and in general, for time-based interactions) it is desirable to have *only one* pre-specified action relevant to the context of viewing. For example, the action can be that of zooming-in to the observed detail of the painting (see Figure 6), or proceeding to the next item in the museum collection. The drawbacks of time-based interaction solutions triggered by focusing on the object/area of interest areas follows:

- the problem of *going back* to observation mode. This means that the action triggered by focusing on a certain area has to be either *self-terminating* (as is the case with the 'display the next artifact' action, where the application switches automatically back to the observation mode) ,or one has to provide a simple mechanism that would allow the viewer to return to the observation mode (for example, by moving the gaze focus outside of the object boundary);
- the problem of choice between *multiple actions*. Using the time-based mechanism, it is possible to trigger different actions. By changing the cursor/ focus shape, one can also indicate to the viewer which action is going to take place. However, since the actions are tied to the objects themselves, the viewer essentially has no choice but to accept the pre-specified action. This may not be a problem in a context where pre-specified actions are meaningful and correspond to the viewer's expectations. However, it does limit the number of actions one can 'pack' into an application and can create confusion in cases where two instances of focusing on the same object may trigger off different actions.
- the problem of *interrupted flow* or *waiting*. Inherent to time-based solutions is the problem that the viewer always has to wait for an action to be executed. In my experience, after getting acquainted with the interaction mechanism, the waiting time becomes subjectively longer (because the users know what to expect) and often leads to frustration. The problem can be diminished to some extent by progressively shortening the duration of focusing necessary to trigger the action. However, at some point it can lead to another

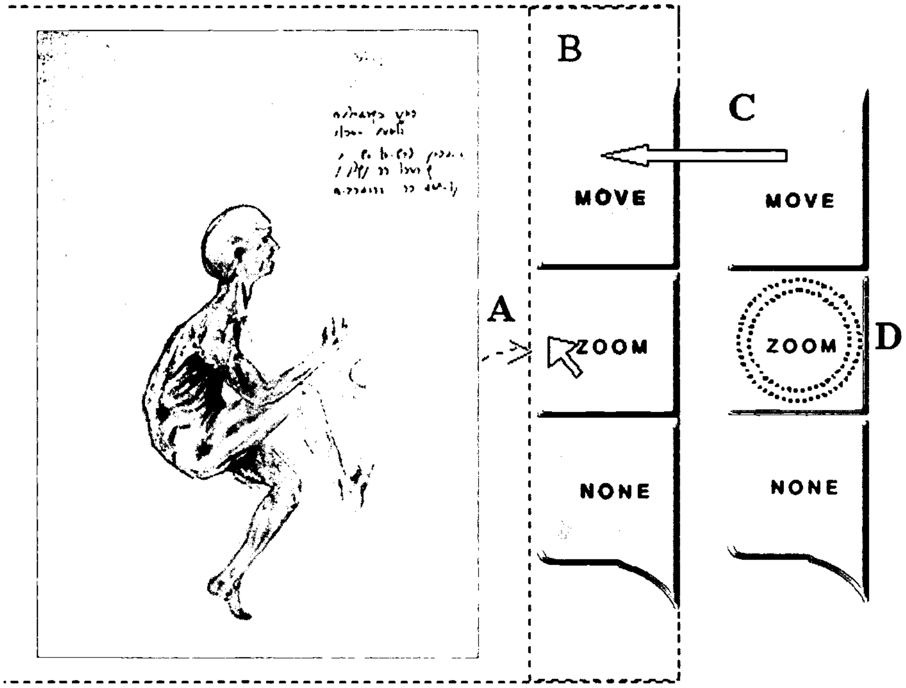

Figure 6. After choosing the desired action (see Figure 5), returning the gaze to the object executes the action without delay. The detail above shows the 'zoom-in' tool, which becomes 'tied' to the viewer's gaze and allows close inspection of the artifact.

source of frustration since the viewer may be forced to constantly shift the gaze around in order to stay in the observation mode.

Inspite of the above mentioned problems, time-based gaze interactions can be an effective solution for museum use where longer observation of an area of interest provides the viewer with more information. Another useful approach is to use the gaze direction as input for the delivery of additional information through another modality. In this case, the viewer does not need to get visual feedback related to his/her eye movements (which can be distracting on its own). Instead, focusing to an area of interest may trigger voice narration related to viewer's interest. For an example of this technique in the creation of a gaze-guided interactive narrative, see Starker & Bolt (1990).

Location-based interfaces

Another traditional way of solving the "clutch" problem in gaze-based interfaces is by separating the modes of observation and action by using controls that are in the proximity of the area of interest but do not interfere with visual inspection. I have already described EyeCons (Figure 3) designed by the Risø National Research Laboratory in Denmark (for a detailed description see Glenstrup and Engell-Nielsen, 1995). In the following section I will first expand on EyeCons design and then propose another location-based interaction mechanism. The first approach is illustrated in Figure 5.

Fixating any of the buttons is equivalent to a button press and chooses the specified action which is executed without delay when the gaze returns to the object of interest. The viewer can also return to observation mode by choosing no action button. The action palette slides out of view as soon as the gaze moves out of the area (B).

The observation area (the drawing) and the controls (buttons) are separated. At first glance, the design seems very similar to that of the EyeCons, but there are some enhancements that make the interactions more efficient. First, the controls (buttons) are located on a configurable 'sliding palette',

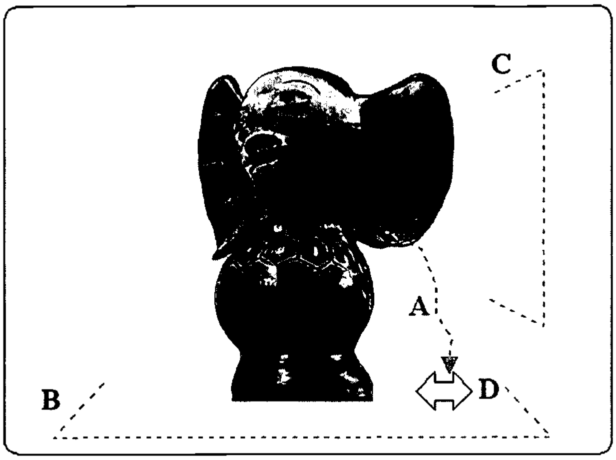

Figure 7. Surface-based interaction mechanism. Viewer's scanpath is visible at (A). Two interactive surfaces (B and C) are discretely marked on the projection. Moving the gaze into the area of interactive surface is marked by appearance of cursor with the shape indicative of possible action (D).

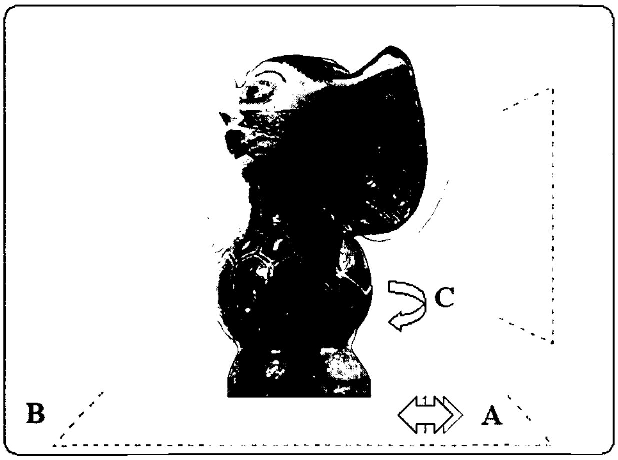

Figure 8. If the viewer's gaze (as indicated by cursor position at A) remains within interactive surface (B), any gaze movement within the surface will lead to incremental action – in this case rotation of the object (C).

a mechanism that was adopted by the most widely used operating system (Windows) in order to provide users with more 'screen real estate'. The reason for doing this in a museum context is also to minimize the level of distraction while observing the artifact. Shifting the gaze to the side of the projection space (B) slides the action palette into the view. The button that is currently focused becomes *immediately* active (D) signaling the change of mode by displaying the focus ring and changing the color. This is a significant difference compared to the EyeCons design, which combines both location- and time-based mechanisms to initiate action. Moving the gaze back to the object leads to the execution of specified action (selection, moving, etc.). Figure 6 illustrates the outcome of choosing the 'zoom' action from the palette. The eye-guided cursor becomes a magnifying glass allowing close inspection of the artifact.

One can conceptually expand location-based interactions by introducing the concept of an *active surface*. Buttons can be viewed as being essentially single-action locations (switches). It really does not matter which part of the button one is focusing on (or physically pressing) – the outcome is always the same. In contrast, a surface affords assigning meaning to a series of locations (fixations) and makes possible *incremental* manipulation of an object.

Figure 7 provides an example of a surface-based interaction mechanism. Interactive surfaces are discretely marked on the area surrounding the object. For the purpose of illustration, a viewer's scan path (A) is shown superimposed over the object and indicates gaze movement towards the interactive surface. Entering the active area is marked by the appearance of a cursor in a shape that is indicative of the possible action (D). The appearance of the cursor is followed by a brief latency period (200-300 ms) during which the viewer can return to the observation mode by moving the gaze outside of the active area. If the focus remains in the active area (see Figure 8), any movement of the cursor along the longest axis of the area will be *incrementally mapped* onto an action sequence – in this case, rotation of the object.

The advantages of surface-based interaction mechanisms are the introduction of more complex, incremental action sequences into eye movement input and the possibility of rapid shifts between the observation and action modes. The drawback is that the number of actions is limited and that the sur-

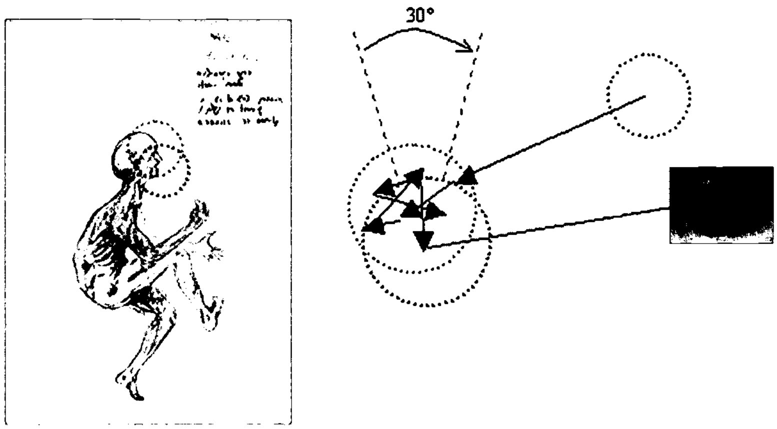

Figure 9. Gaze-pointing. The viewer can observe the artifact with the pointing cursor (dashed circle) indicating the current gaze location. Sweeping gazes across the scene are possible as long as they are not in upward direction and end in the 30° angle strip.

faces, although visually non-intrusive, still claim a substantial portion of the display.

Action-based interfaces

Building on the previous two models, one can further expand the conceptual framework for gaze-based interfaces. This time I will focus on the *gaze action* as a mechanism for switching between the observation and the active (command) mode. Analysis of the previously described surface-based model reveals that it can be described as an intermediary step between the surface- and action-based interfaces. In this model, although the shift between the observation and action mode is dependent on *the location* of gaze focus, the control of interaction is based on gaze action (moving the focus/cursor over gaze-sensitive surface). Thus, the last step in our analysis is to explore the possibility of using predominantly *gaze-based actions* as a control mechanism. This may seem like slippery ground because physiologically our visual behavior is mostly geared towards collecting information and not acting upon the world. The exception of a kind is in the domain of sexual and social behaviors where gaze direction and duration may literally have physical consequences by signaling attraction, dominance, submissiveness, etc. Fine literature abounds with examples describing gazes as having a tangible effect ("his piercing gaze," "he felt her gaze boring two little holes at the back of his neck...," "her angry gaze was whipping across the room trying to find out who did this to her.." to mention a few). Our ability to transfer knowledge from one sensory domain to another modality will be the key component in the proposed outline of an action-based gaze interface.

In eye-tracking literature, a gaze is most often defined as a number of consecutive fixations in a certain area. This metric emphasizes the *location* and the *duration* characteristics of the gaze and can be extremely useful in inferring the viewer's interest or gauging the complexity of the stimulus. However, in my proposal I would like to focus on two, often neglected, characteristics of a *moving gaze* that can be consciously used by a viewer to indicate his/her *intention*. These are:

- the direction of gaze movement, and
- the speed of gaze movement.

For technical purposes a moving gaze can be defined as a number of consecutive fixations progressing in the same direction. It corresponds roughly to longer, straight parts of a scan path and is occasionally referred to as a *sweep* (Altonen et al. 1998). The reason for choosing these characteristics is twofold. First, eyes can move much faster than the hand (and there is evidence from literature that eye-pointing is significantly faster than mouse pointing, see Sibert and Jacob 2000). Second, as mentioned before, directional gaze movement is often used in social communication. For example, we often indicate in a conversation exactly 'who' we are talking

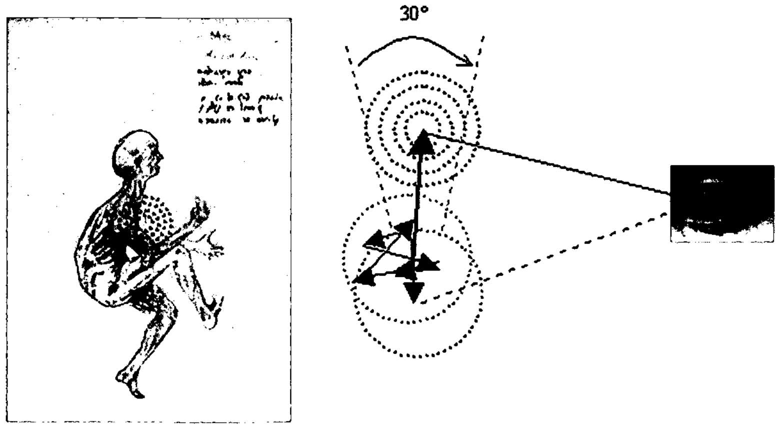

Figure 10. Gaze-selection. Shifting the gaze rapidly upwards within the 30° triggers of the selection process. The cursor changes the shape to that of a target and positions itself at the center of the object as a prelude to the action of gaze-dragging.

about by repeatedly shifting the gaze in the direction of the person in question.

In order to create an efficient gaze-based interface, one has to be able to replicate the basic mouse-based actions used in the traditional graphical user interface (GUI). These are: **pointing** (cursor over), **selection** (mouse down), **dragging** (mouse down + move) and **dropping** (mouse up). I will also propose the inclusion of yet another non-traditional action, which I introduced in interface design a while ago (Milekic, 2000) and which proved to work extremely well as an intuitive browsing mechanism. This is the action of *throwing* which is dependent on the speed of movement of a selected object. Compared to the traditional interface, the throwing action is an expansion of the action of dragging an object. As long as the speed of dragging remains within a certain limit, one can move an object anywhere on the screen and drop it at desired location. However, if one 'flicks' the object in any direction, the object is released and literally 'flies away' (most often, to be replaced by another object). I have implemented this mechanism in a variety of mouse-, touchscreen- and gesture-based installations in museums and it has been successfully used by widely diverse audiences, including very young children. Subjectively, the action is very intuitive and natural, and the feeling can be best compared to that of sliding a glass on a polished surface (a skill that many bar tenders hone to perfection). In the following sections I will describe each of the gaze-based actions.

Gaze-pointing (Figure 9) is the easiest function to replicate in a gaze-based interface. It essentially consists of a visual clue that indicates to the viewer which area of the display is currently observed. Although one can use the traditional cursor for this purpose, it is desirable to design a cursor that will not interfere with observation. Dynamic change of cursor shape when moving over different objects can also be used to indicate whether an object is gaze-sensitive and to specify the type of action one can initiate (this technique is used in surface-based interface, described above; see Figure 4, for example). I have chosen a simple dashed circle as an indicator of the current gaze location. Pointing action is maintained as long as there are no sudden substantial changes in a specific gaze direction. If such a change occurs, the tracking algorithm determines the *direction* of gaze movement and, if necessary, initiates appropriate action.

This does not mean that the viewer is limited to slow (and unnatural) observation. In fact, switching from observation to action mode (selection) occurs *only* if movement of sufficient amplitude occurs in an upward direction and ends up in a fairly narrow area spanning approximately 30° above the current focus area. This means that viewers can, more or less, maintain a normal observation pattern, even if it includes sweeping gaze shifts, as long as they don't end up in the critical area.

Gaze-selection (Figure 10) is an action initiated by a sudden upward gaze shift. The action is best de-

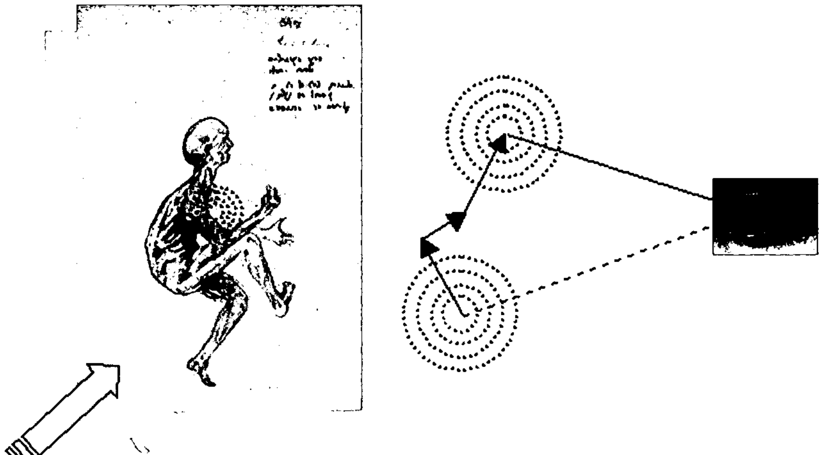

Figure 11. Gaze-dragging. The painting is 'hooked' to viewer's gaze and follows its direction. At this stage the viewer can decide either to 'drop' the painting at another location (see Figure 12) or, 'throw' away the current one and get a new artifact.

scribed (and subjectively feels like) the act of upward stabbing, or 'hooking' of the object. In a mouse-based interface the selection is a separate action – that is, one can just select an object, or select-drag-drop it somewhere else, or de-select it. In a gaze-based interface, what happens after the selection of an object will depend on the context of viewing. When multiple objects are displayed, the selection mechanism can act as a self-terminating action, making it possible for the viewer to select a subset of objects. In this case, highlighting the object would indicate the selection. However, in the museum context (assuming that the viewers will most often engage in observation of a single artifact) object selection may just be a prelude to the action of moving (dragging). In this case the object becomes, figuratively speaking, 'hooked' to the end of the viewer's gaze, as indicated by a change of the cursor's shape to that of a target.

Gaze-dragging (Figure 12). Once the object has been selected ('hooked' to the viewer's gaze), it will follow the viewer's gaze until it is 'dropped' at another location. This action is meaningful in cases when the activity involves the repositioning of multiple objects (for example, assembling a puzzle). In the scenario depicted above, the viewer can 'throw away' the current object and get a new one.

Gaze-throwing (Figure 13) is a new interaction mechanism that allows efficient browsing of visual data bases with a variety of input devices, including

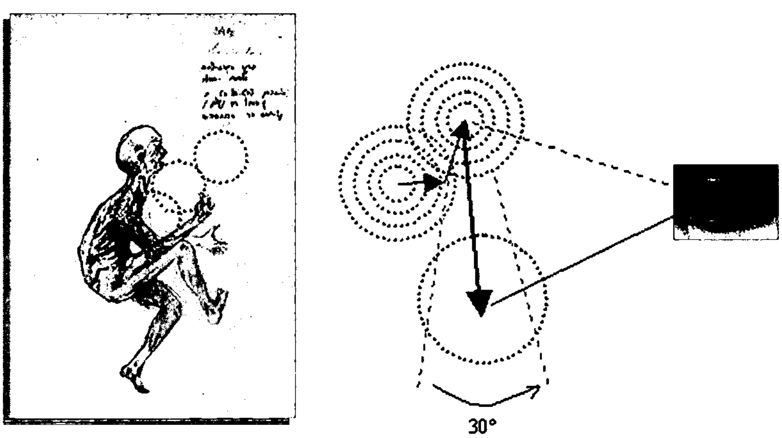

Figure 12. Gaze-dropping. The action of dropping an object is the opposite of 'hooking' it. A quick downward gaze movement releases the object and switches the application into observation mode.

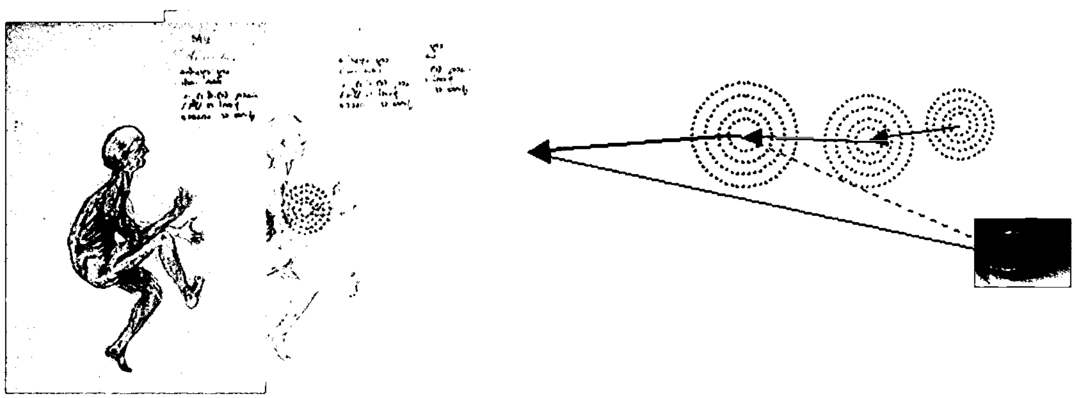

Figure 13. Gaze-throwing. 'Throwing' an object away is accomplished by moving the gaze rapidly to the left or to the right. Once the object reaches threshold speed it is released and 'flies away'. A new artifact floats to the center of display.

gaze input. An object that has been previously selected ("hooked") will follow the viewer's gaze as long as the speed of movement does not exceed a certain threshold. A quick glance to the left or the right will release the object and it will 'fly away' from the display to be replaced by a new artifact.

The objects appear in a sequential order, so if a viewer accidentally throws an object away, it can be recovered by throwing the next object in the opposite direction.

To summarize, action-based gaze input mechanisms have the advantage of allowing the viewer to *act upon the object* at will, without time or location constraints. The mechanism is simple and intuitive be-

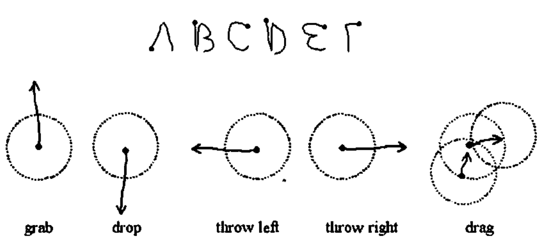

Figure 14. Eye-graffiti. Top row presents graffiti used for text input (letters A,B,C,D,E,F respectively) in Palm OS based personal organizers. Bottom row outlines suggested gaze-gestures that trigger different actions once the object has been selected.

cause it is analogous to natural actions in other modalities. The best way to think about action-based gaze input is as a kind of *eye-graffiti*. The vocabulary of suggested gaze-gestures for eye input is presented in Figure 14. It is similar to the text input mechanism used for Palm personal organizers where the letters of the alphabet are reduced to corresponding simplified gestures. The fact that millions of users were able to adopt this quick and efficient text input mechanism is an indication that the development of *eye-graffiti* has significant potential for gaze based interfaces.

The dashed circle in the illustration does not represent the visual representation of the cursor, but rather the area used to calculate the direction and the velocity of gaze movement by the tracking algorithm. The heavy dot indicates the starting point of a gesture. However, while action-based gaze input mechanism may seem best suited for museum applications, the ideal interface is probably a measured combination of all three approaches.

III. Current Technologies for Non-Intrusive Eye Tracking

Unlike in the laboratory environments, the eye-tracking technology used in a museum setting has to meet additional specific requirements. Some of the most obvious ones are:

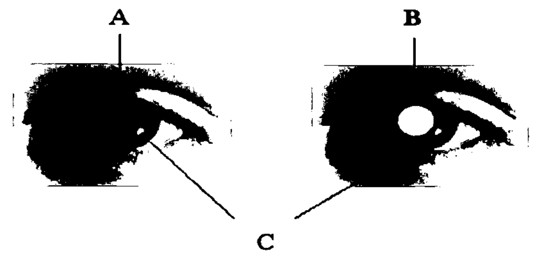

Figure 15. Gaze direction can be calculated by comparing the relative position and the relationship between the pupil (A) and corneal reflection – the glint (C). Infra-red illumination of the eye produces the 'bright pupil' effect (B) and makes the tracking easier.

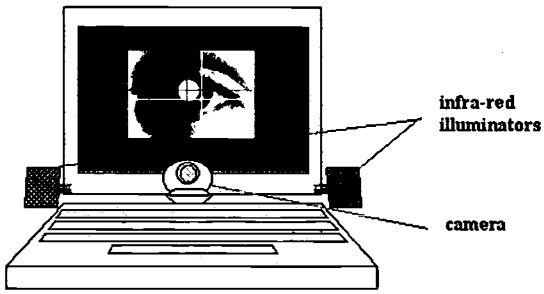

Figure 16. Several manufacturers produce portable eye-tracking systems similar to the one depicted above. While the camera position is most often bellow the eye level (eyelids interfere with tracking from above), the shape and position of infrared illuminators vary from manufacturer to manufacturer.

- it should be non-intrusive. This excludes all eye-tracking devices that use goggles, head-straps, chin-rests or such.

- it should allow natural head movements that occur during viewing.

- it should not require individual calibration.

- it should be able to perform with a wide variety of eye shapes, contact lenses or glasses.

- it should be portable.

- it should be affordable.

With the ncreasing processor speeds of currently available personal computers, it seems that the most promising eye-tracking technology is that based on digital video analysis of eye movements. The most commonly used approach in video-based eye tracking is to calculate the angle of the visual axis (and the location of the fixation point on the display surface) by tracking the relative position of the pupil and a speck of light reflected from the cornea, technically known as the "glint" (see Figure 15). The accuracy of the system can be further enhanced by illuminating the eye(s) with low-level infra-red lightto produce the "bright pupil" effect and make the video image easier to process (B in Figure 15). Infrared light is harmless and invisible to the user.

A typical and portable eye-tracking system similar to the ones commercially available is depicted in Figure 16. Since the purpose of this paper is not to endorse any particular manufacturer, I urge interested readers to consult the large Eye Movement Equipment Database (EMED) available on the World Wide Web. Keeping in mind that many museums and galleries have very modest budgets, I will specifically address the issue of *affordable* eye-tracking systems.

The price range of most commercially available eye-trackers is between $5000 and $60.000, often with additional costs for custom software development, setup etc. Although there are some exceptions, the quality and the precision of the system tend to correlate with the price. However, with the increasing speed of computer processors, greater availability of cheap digital video cameras (like the ones used for Web-based video conferencing) and, most importantly, the development of sophisticated software for video signal analysis, it is becoming possible to build eye-trackers within a price range comparable to that of a new personal computer. Even though the cheaper systems have lower spatial and temporal resolution when compared to the research equipment, in a museum/gallery setting they may be used for different applications; for example, for browsing a museum collection with additional information provided by voice-overs. A more significant use would be providing access to the museum content to visitors with special needs. An example of a cost-effective solution based on a personal computer and a Web-cam for eye-gaze assistive technology was recently described (Corno, Farinetti and Signorile, 2002).

Most commercially available eye-tracking systems (including the high-end ones) have two characteristics that make them less than ideal for use in museums. These are:

- the system has to be calibrated for each individual user
- even remote eye-trackers have very low tolerance for head movements and require the viewer to hold the head unnaturally still, or to use external support like head- or chin-rests.

The solution lies in the development of software able to perform eye-tracking data analysis in more natural viewing circumstances. A recent report by Quiang and Zhiwei (2002) seems to be a step in the right direction. Instead of using conventional approaches to gaze calibration, they introduced a procedure based on neural networks that incorporates natural head movements into gaze estimation and eliminates the need for individual calibration.

The emergence of eye-tracking technologies based on a personal computer equipped with a Web-cam and the development of software that allows gaze tracking in natural circumstances open up a whole new area for museum applications. The described technologies make Web-based delivery of gaze-sensitive applications possible. This not only presents an opportunity for a novel method of content delivery (and reaching different groups of users with special needs) but also offers an incredible possibility to collect, on a massive scale, data related to visual analysis of museum artifacts. However, a word of caution is in order here. One cannot overemphasize the importance of *context* in an eye-tracking application (or, for that matter, in any application). In an appropriate context, even a fairly simple setup can produce magical results, and the use of the most expensive equipment can lead to viewer frustration in a flawed application.

Conclusion

I have outlined a conceptual framework for the development of a gaze-based interface for use in a museum context. The major component of this interface is the introduction of ***gaze gestures*** as a mechanism for performing intentional actions on observed objects. In conjunction, an overview of suitable eye-tracking technologies was presented with an emphasis on low cost solutions. The proposed mechanism allows the development of novel and creative ways for content delivery both in a museum setting and via the World Wide Web. An important benefit of this approach is that it makes museum *content* (and not just the building or the restrooms) accessible to a wide variety of populations with special needs. It also offers the possibility of data-logging related to visual observation on a massive scale. These records can be used to further refine the content delivery mechanism and to promote our understanding of both the psychological and the neurophysiological underpinnings of our relationship with the Art.

References

Altonen, A., A. Hyrskykari, K. Raiha. (1998). 101 Spots, or how do users read menus? in Proceedings of CHI 98 Human Factors in Computing Systems, ACM Press, pp 132-139.

Bever, T.G., (1970). The cognitive basis for linguistic structure, in J.R. Hayes, ed., *Cognitive development of language*, Wiley, New York.

Card, S.K. (1984). Visual search of computer command menus, in H. Bouma, D.G. Bouwhuis (Eds.) *Attention and Performance X, Control of Language Processes*, Hillsdale, NJ, LEA.

Corno, F., L. Farinetti, I. Signorile (2002). A Cost-Effective Solution for Eye-Gaze Assistive Technology, ICME 2002: IEEE International Conference on Multimedia and Expo, Lausanne, Switzerland.

Corkum V., Moore C. (1998). The origins of joint visual attention in infants. Developmental Psychology, 34, pp 28–38.

Dodge, R., T.S. Cline (1901). The angle velocity of eye-movements. Psychological Review, 8, 145–57.

Duchowski, A.T. (2002). Eye Tracking Methodology: Theory and Practice, Springer Verlag.

Glenn III, F.A., H.P. Iavecchia, L.V. Ross, J.M. Stokes, W.J. Weiland, D. Weiss, A.L. Zakland, (1986). Eye-voice-controlled interface, Proceedings of the Human Factors Society, 322-326.

Glenstrup, A.J., T. Engell-Nielsen, (1995). Eye Controlled Media: Present and Future State. Minor Subject Thesis, DIKU, University of Copenhagen, available at: http://www.diku.dk/~panic/eyegaze/article.html#contents.

Hartridge, H., L.C. Thompson, (1948). Methods of investigating eye movements, British Journal of Ophthalmology, 32, pp 581-591.

Hutchinson, T.E., K.P. White, W.N. Martin, K.C. Reichert, L.A. Frey, (1989). Human-Computer Interaction Using Eye-Gaze Input, IEEE Transactions on Systems, Man, and Cybernetics, 19, pp 1527-1534.

Jacob, R. J. K. (1993). Eye-movement-based human-computer interaction techniques: Toward non-command interfaces, in H. R. Hartson & D. Hix, (eds.) *Advances in Human-Computer Interaction*, Vol. 4, pp 151-190, Ablex Publishing Corporation, Norwood, New Jersey.

Jacob, R.J.K., K.S. Karn, (2003). Eye Tracking in Human-Computer Interaction and Usability Research: Ready to Deliver the Promises (Section Commentary), in *The Mind's Eyes: Cognitive and Applied Aspects of Eye Movements*, J. Hyona, R. Radach, H. Deubel (Eds.), Oxford, Elsevier Science (in press).

Judd, C.H., C.N. McAllister, W.M. Steel, (1905). General introduction to a series of studies of eye movements by means of kinetoscopic photographs, in J.M. Baldwin, H.C. Warren & C.H. Judd (Eds.) Psychological Review, Monograph Supplements, 7, pp1-16, The Review Publishing Company, Baltimore.

Just, M.A., P.A. Carpenter, (1976). The role of eye-fixation research in cognitive psychology, Behavior Research Methods & Instrumentation, 8, pp 139-143.

Levine, J.L. (1981) An Eye-Controlled Computer, Research Report RC-8857, IBM Thomas J. Watson Research Center, Yorktown Heights, N.Y.

Milekic, S. (2000) Designing Digital Environments for Art Education / Exploration, Journal of the American Society for Information Science, Vol. 51-1, 49-56.

National Gallery, London "Telling Time Exhibition", Web source consulted 1/24/03, available at: http://ibs.derby.ac.uk/gallery/updates.shtml

Scaife M., J. S. Bruner (1975). The capacity for joint visual attention in the infant. Nature, 253, pp 265–266.

Sibert, L. E., R.J.K. Jacob (2000). Evaluation of Eye Gaze Interaction, Proceedings of the CHI 2000, ACM, New York, 281-288, available at: http://citeseer.nj.nec.com/article/sibert00evaluation.html

Starker, I., R.A. Bolt, (1990). A gaze-responsive self-disclosing display, in *CHI '90 Proceedings*, ACM, pp. 3-9.

Stewart, J., & C. Logan, (1998). *Together: Communicating Interpersonally (5th Ed.)*. Boston: McGraw Hill, 90-100.

Qiang, J., Z. Zhiwei, (2002). Eye and Gaze Tracking for Interactive Graphic Display, International Symposium on Smart Graphics, June 11-13, 2002, Hawthorne, NY

Vertegaal, R. (1999). The GAZE groupware system: mediating joint attention in multiparty communication and collaboration, in *Proceedings of the ACM CHI'99 Human Factors in Computing Systems*, ACM Press, New York, pp 294-301.

Ware, C., H.H. Mikaelian, (1987). An evaluation of an eye tracker as a device for computer input, Proceedings of the CHI 1987, ACM, New York, 183-188.

Wiley.Moore, C. (1999) Gaze following and the control of attention, in P. Rochat (ed.), *Early social cognition: understanding others in the first months of life*, pp 241–256. LEA, New York

Yarbus, A.L. (1967) Eye movements during perception of complex objects, in L.A. Riggs, (ed.), *Eye Movements and Vision*, Plenum Press, New York, 171-196.

Zhai, S., C. Morimoto, S. Ihde, (1999). Manual and Gaze Input Cascaded (MAGIC) Pointing, Proceedings of the CHI'87, ACM New York, 246-253.

Re-assessing Practice: visual art, visually impaired people and the Web

Caro Howell, Tate Modern and Dan Porter, Freelance Web Developer/Art Historian, United Kingdom

http://www.tate.org.uk/imap/

Abstract

Interpretation for visually impaired people in art museums is dominated by an emphasis on the tactile and the physically immediate. In practice, this seemingly logical approach often keeps blind and partially sighted people at an intellectual distance from art works and their artistic context. For over-reliance on a tactile approach has the effect of making sculpture the primary vehicle for accessing art, despite the fact that "touchable" sculpture represents only a tiny fraction of Western art as represented in museums and galleries. Moreover, in this display context "touchable", by definition, means robust, indicating traditional materials such as bronze and stone. Yet like traditional narrative and figurative subject matter, bronze and stone cease to hold sway in twentieth century art. How then can someone who is congenitally blind be given intellectual access to non-tactile artworks that are not artefacts, that do not have clear descriptive relationships to objects and experiences from the lived world and that refer to and are of an about themselves only?

This was the starting point for research that began four years ago at Tate Modern, London's new national museum of modern art. The latest development to come out of this on-going research is i-Map art resources for blind and partially sighted people that are delivered on-line. Currently i-Map explores the work of Matisse and Picasso, their innovations, influences and personal motivations, as well as key concepts in modern art. Aimed at partially sighted and blind people with a general interest in art as well as art teachers and their visually impaired pupils, i-Map incorporates text, image enhancement and deconstruction, animation and raised images. Importantly, i-Map transformed a gallery-based practice that involved intensive 1:1 delivery, into an entirely new way of deconstructing art on-line and one where the user has the necessary tools to work independently. Moreover, i-Map goes beyond straight description, attempts to simulate purely visual experiences and the usual focus on "what?" in favour of exploring the "why?" of art so that visually impaired users can make their own critical judgments.

Visually impaired people use the Web in conjunction with screen reader software to obtain information, visually impaired people often find travel complicated and stressful, visually impaired children are usually in mainstream education and all schools in the UK are on-line. The process of reassessing the parameters and definitions of art education for visually impaired people revealed the Web to be an ideal vehicle for increasing intellectual access and delivering effective interpretation in a format that offers autonomy of exploration. However, in order for a project such as i-Map to confidently defy received Web design wisdom and develop tailor made solutions, it's content needs to be the product of successful methodology and focussed research. It is possible to provide blind and partially sighed people with intellectual access to any artwork and the Web offers enormous potential to do assist in this process. In attempting to achieve this, i-Map can provide useful practical and pedagogical pointers.

'I was born blind, as a child I had a little sight, enough to see light and dark and a few colours, though not enough to see pictures or perceive any kind of object other than its mass. Therefore, the concept of visual art is quite new to me and, although I've been through some of these descriptions a couple of times I don't think I've grasped some of the finer points as some of the ideas strike me as very visual. I hope to be able to appreciate these concepts more in future by thinking about them and talking to other people.'

This was a comment emailed to us via i-Map's feedback form. i-Map is an on-line project designed to give blind and partially sighted people access to the work of Matisse and Picasso. What is striking about this person's comment is the familiarity of the response. How many of us have had exactly that reaction when looking at art, especially modern art? The feeling that you're missing the point, but the sense that thinking it over and talking to others will clarify things. For Caro Howell and Dan Porter , the creators i-Map, this comment (and others like it) also suggests that our ambitions for the project have been realised, namely that a blind person has independently engaged with and thought about visual art without having to physically touch it.

i-Map came about as a result of work done by Caro Howell at Tate Modern with visually impaired people. Tate Modern is the UK's national museum of modern art. It opened in 2000 and Caro was part of a small team who in 1997 began developing a pre-opening public programme for Tate Modern and post-opening exhibitions, displays, education and interpretation. Part of her brief was to consider issues of physical and intellectual access. From experience she knew that traditional gallery education programmes for visually impaired audiences were not going to succeed at the new museum since they focused on touch, either through the use of dedicated handling collections, or, more often, through the supervised handling of artworks considered suitable by conservation staff.

Tate Modern shows art from 1900 to the present day, so the number of robust sculptures suitable for handling and on display at any one time is very low. Moreover, the gallery is huge – the building is a former power station – and displays are shown over three floors. So cherry picking only 'touchable' artworks would turn a visit into a forced march. Of primary importance though, was the issue of intellectual access. How were visually impaired people to engage with and understand the radical developments that have taken place in Western art since 1900, if they were only ever to experience three dimensional works and then only in bronze and stone?

Project History

In 1999 Caro began working with a group of visually impaired adults who had different degrees of impairment and different experiences of art, to explore language as an interpretative tool. In addition to guided tours, Tate Modern was planning self-directed audio tours for visually impaired visitors and Caro was interested in exploring different approaches to audio description. Traditional methods used in the theatre are neutral and objective. This approach is suitable if audio description merely fills in the contextualising gaps around the actors' dialogue, but it's clearly unsuitable for describing even figurative art.

The group worked with sculptures that could be explored through touch, but with the aim of creating a blueprint for interpreting art that couldn't be touched. Since the group's day-to-day experience of touch was pragmatic (as it is for most visually impaired people), initially their automatic descriptive responses to tactile sensations tended to be generalised and unimaginative – warm, rough, big etc. So creative writers and performance poets were brought in to help them experiment with different structures and methods of description. These writers enabled the group to explore language as they explored different tactile sensations, to think, feel and articulate creatively. The challenge that emerged was how to describe with accurate particularity. In other words, identify quintessential qualities in an artwork and articulate them in ways that conveyed their uniqueness, but also made sense to others. It was this process that revealed the power of simile and metaphor as tools for enabling a group of people with disparate visual and personal experiences to reach quickly a shared point of understanding about something they couldn't see. However, metaphors and similes are also tricky to control. Their power to tap into people's imagination is itself a liability, because if not chosen correctly, they take the mind off on erroneous tangents that run counter to the artist's intentions. For instance, describing the surface of a sculpture as like candyfloss – cotton candy – when the piece is about mankind's survival post World War II, European occupation and the holocaust is clearly not appropriate, even though in many respects the suggested candyfloss similie was logical.

So, rigorous thinking about language was one important result of the Making Sense project and this would provide useful guidance when it came to writing i-Map. Another revealing discovering was the ability of tactile stimuli to provide physical metaphors for an artwork's effect or an artist's intention. Sometimes, tactile sensations were found to be more revealingly than language in capturing the essence of an artwork. An example of this that always springs to mind is a silicone breast implant that is a handling object used to explore a particular painting by Salvador Dali. The implant manages to capture the painting's disturbing combination of attraction/repulsion, familiarity/surreality and in particular, the quality of the painting's central feature which is an entropic hunk of flesh that seems to simultaneously seep and yet hold its shape.

The linguistic and tactile findings of the Making Sense project led directly to the formulation of Tate Modern's touch tours and, in turn, the methodology and philosophy of the tours informed i-Map. The Collection displays at Tate Modern are displayed not chronologically, but thematically across four large suites of rooms. The themes are based on the traditional genre of art – history painting, the nude, landscape and still life. Any visitor to the gallery has to engage with this curatorial intervention, even if they engage with nothing else and Caro felt strongly that the same should be true for visually impaired visitors. So four touch tours were created, one for each of the suites. Each tour is designed to explore the theme as a whole, as well as engage with a particular aspect of that theme. For instance, the tour in the Still Life/Object/Real life suite explores still life painting as a genre, in addition to looking at the way artists' use of materials has changed since 1900. When this tour was created is was considered by many to be the most challenging, as it contained almost not art that could be touched and included work by Picasso and Duchamp, artists who had previously been considered too 'difficult' to study with visually impaired people. However, this tour has proved to be the most successful with visitors precisely because of its challenging content. Far from being obscure and meaningless, blind and partially sighted people find the visual complexity of Cubism, collage and ready-mades and the artists' engagement with the conventions of art fascinating. As one blind woman put it at the end of the tour, having at first stated that she didn't like modern art; 'This makes sense to me because I think conceptually too.' At the heart of Tate Modern's and i-Map's approach to working with visually impaired people is the belief that engaging with art is not simply a question of looking at it. If it was that simple, people wouldn't have a problem with modern art, many of us would be out of a job and we'd all switch off the radio or put down the phone every time someone started talking about something we couldn't see. Looking at art is a process that engages the intellect, the emotions, personal lived experience as well as a phenomenological sense of self in relation to the object. Of course ocular acuity is an integral part of engaging with visual art and lacking sight does have a profound impact on one's relationship to art, but blindness alone doesn't negate or preclude engagement. In fact the notion of the ideal eye is itself a flawed one. It is said that one in 20 men are colour blind, Monet was painting with cataracts and all those people who don't wear contact lenses or glasses, when was the last time you had your sight tested? Are you sure you're seeing perfectly? In reality, a substantial proportion of visitors to our museums and galleries must, by the law of averages, not be seeing in physical terms what the artist saw, let alone the mediating curators and interpreters. So there's a case to argue that if visually impaired people can't look at art in public museums and galleries the fault lies with us, museum professionals, and our lack of commitment, imagination and ambition, not with the visitor. The the word 'look' is used deliberately, as it's used by visually impaired people, to mean to engage with and to understand.

The touch tours had revealed that no art was out of bounds, be it painting or video; installation or sculpture; figurative or abstract. More importantly, the traditional fixation with touch – 'if we can't touch it we can't use it' – was successfully challenged. After all, what would be the point of touching Duchamp's 'Fountain', which is a urinal when it's not the original urinal (that got lost and this is one of fifteen replicas), it's not in the original porcelain, but glazed earthenware painted to look like porcelain and more importantly, the work isn't about the urinal, but the nature of art and the role of the artist. Nevertheless, the numbers of visually impaired people visiting Tate Modern and having touch tours is small, perhaps one hundred and fifty a year. This is due in no small part to the logistical difficulties blind and partially sighted people face when planning a visit. Moreover, provision in art museums for visually impaired people has historically been poor, so there's not the expectation that a visit will be rewarding. So it takes a remarkable level of inquisitiveness for a visually impaired person with no previous relationship to art to decide to visit Tate Modern, on the off chance that contemporary art will be their thing. The challenge was therefore to find a way to bring the pedagogy of the touch tours to a wider visually impaired audience who may never come to the gallery. One group that was identified was visually impaired children. In the UK the majority of visually impaired children, if they have no other disabilities are in mainstream education. Many of them decide to study GCSE and A-Level art alongside their sighted classmates. However, they are taught by art teachers who have often little or no specialist training in teaching the curriculum to students with little or no useful sight. As a result, many visually impaired children end up being sidelined because they're set projects that bear no relation to the work their peers are undertaking, or that engage critically with the history of art. A graphic example of this was a GCSE class that was brought to the gallery by their teacher. The students were all coming to see examples of artists who use strong colour, as they were studying Fauvism. All that is except the one visually impaired student. She had lost her sight at the age of seven and had limited visual memory, so she was doing a project about dragons. The attitude towards this pupil was very much 'can you entertain her whilst the rest of the class studies the art.' It had clearly never occurred to the teacher that her pupil could be included in discussion about Fauvism or that she, like most visually impaired people including congenitally blind people have a sophisticated conceptual sense of colour.

Naturally, this lack of involvement at an early age will often lead to a complete alienation from the world of visual art. Perhaps the most telling comment we heard when testing i-Map came from a blind man who, after struggling to extract some meaning from a description of Matisse's *Moroccans* said, 'to be honest, I'm just not interested in art, I guess I've always thought it just wasn't for me.' Indeed, he seemed confused as to why we were even pursuing the project. Given that he had always felt vaguely excluded from art, and thus had no contextual framework through which to interpret it, this lack of engagement should not have come as a surprise. Yet the touch tours at Tate had proved that elsewhere, there were visually impaired people with a hunger to learn more. As Caro had found, these individuals often presented a different challenge. Keen to develop ways of understanding art for themselves, they would often arrive with ill informed or half-understood notions, overheard from sighted descriptions and knitted together with uncertainty. On one hand, expert accounts of artists and styles made no allowances for visual impairment, while on the other, the accounts of helpers, friends and family were usually non-expert. Either way, outside of specialist schemes such as touch tours, curiosity was most often repaid with a confused and distorted perspective on art. To help break this cycle, we were committed to producing informed professional explanations in deliberately inclusive, non-visual terms. But in addition to the visually impaired users themselves, i-Map clearly needed to educate those whose descriptions and definitions were so often relied upon.

The Development of i-Map

To summarise then, the concept of i-Map was developed in the knowledge that a considered linguistic approach, complimented by ancillary tactile materials, could and had made conceptually and visually complex art accessible. Yet for i-Map, the touch tours were the starting point rather than the model. Touch tour users were rare cases, and arguably motivated as much out of curiosity for the touch tour process as for the art itself. As an on-line project, i-Map could reach a far wider audience. It would be easily available at home and in schools, removing the logistical burden of a gallery visit. And it would bypass the censorship of sighted people presuming visual art to be off limits to the visually impaired. While touch tours relied on interested parties to book an appointment, i-Map, it could be hoped, would awaken in visually impaired Web users an interest in art where none had previously existed. It was in no way a replacement for the touch tours. Indeed, as a Web site it would inevitably lack some of the most powerful elements of the Touch tour methodology; the dialogue, the instant response to questions, the guiding hand of a specialist to keep the inquiry focused and so on. i-Map would have to generate an entirely new set of learning tools, albeit tools with the pedagogical breakthroughs of the touch tours inherent in their design.

The *Matisse Picasso* exhibition at Tate Modern offered a perfect launch pad for a pilot project. Not only did the show have huge, worldwide public interest, but it also featured two artists who had traditionally been both physically and conceptually out of bounds to visually impaired people. In promising to stimulate an engagement with Matisse and Picasso across the entire spectrum of visual impairment we were, in effect, promising intellectual access to concepts such as modernism and cubism, not only to the user who has limited useful sight but also to the user who has never seen: a user, potentially, who has no visual perception of colour whatsoever and no understanding how three-dimensionality might be represented on a flat surface. So even to suppose that it was merely Matisse and Picasso that needed elucidating was to jump the gun considerably. Many of our projected users would need an understanding of picture-making, let alone these relatively late developments in its evolution.

We began by identifying themes that would serve to unite pairs of artworks, one by each artist. Like the exhibition itself, these pairings would allow us to explore both the commonalities and the differences between the work of Matisse and Picasso. But under the generous umbrella headings of 'Primitivism', 'Space' and 'the Nude' we could also begin to build up a broad sense of the techniques, traditions and ideas that are crucial to the interpretation of any modern artwork, and especially so when providing the user with a frame of reference from scratch. To these three headings we would also add 'Cut-Outs.' Though perhaps not a classic genre, this provided a good opportunity for the visually impaired user to identify with some of the more tactile, spatially expressive works, as well as being an elegant synopsis of ideas expressed in more intricate ways elsewhere. Wherever possible we tried to select paintings and sculptures owned by Tate, in order that visually impaired people in the UK had a better chance of visiting them. This would also go some way towards alleviating the copyright difficulties we faced.

The structure and functionality of i-Map had still only been explored tentatively at this stage, but these thematic pairings did at least provide a backbone to the project. On-line visual art projects aimed at the visually impaired were virtually unprecedented, though we were offered a glimpse at a possible way forward by a Web site created by the Finnish National Gallery which explored the work of symbolist painter Hugo Simberg (http://www.fng.fi/fng/html4/en/peda/project/simberg/intro.htm). The method they had used was to isolate and describe in turn various sections of a painting. By studying a section at a time, the visually impaired user was never overwhelmed by detail as they would have been if confronting the whole. They were able to accumulate detail at their own pace and arrive in their own time at an understanding of the complete image. In accessing these descriptions the user was offered a choice of text and audio. Obviously text, easily enlarged or accessed with screen reader software, was a prerequisite of I-Map too. Audio was also discussed, and would indeed have provided a welcome alternative to the expressionless, synthesised drone of the screen reader. Unfortunately, budgetary restrictions eventually prevented us from including it, but it is very much planned as part of our next wave of iterations.

This Finnish example however, could only take us so far. Simberg's works are figurative and strongly narrative, so describing them was a relatively straightforward task. i-Map, on the other hand, was aiming to provide a critical engagement with works that were stylistically radical, non-narrative and requiring an understanding of art's conventions in order to discover their importance. If the work of Matisse and Picasso were to remain legible to the mind's eye, we would need to build into i-Map an ability to communicate the conceptual as well as the concrete, a strategy for text and image which moved beyond the literal descriptiveness of the Simberg site.

Drawing upon the touch tours, Caro was aware that one of the most successful methods of communicating complex pictorial concepts was the use of tactile diagrams, or raised images. These are essentially simple line drawings in which the line has been raised, by means of specialised equipment, to stand in relief against the surface in order to be touch-read. Caro had been developing sets of raised images in which components of the painting were isolated, as stand alone tactile objects, to be interpreted one at a time. Generally, once each element had been explored and understood, the user would be introduced to a raised image that brought these now familiar forms together, representing the artwork as a whole. Given the intentional contradictions and tactics of disorientation characteristic of much modern art, this final step could place heavy demands on the concentration. And yet as a challenging, often-ambiguous interaction it could, with careful guidance, become an equivalent to a sighted engagement with the artwork, and furthermore one in which certain key conceptual milestones would necessarily have been passed.

i-Map would make use of this approach in two ways. First, we would include sets of line drawings available as downloads from the site. Once printed and processed as raised images they could be used in conjunction with the text content of i-Map, extending beyond the hermetic confines of the computer screen to create a further tactile dimension. Secondly, the step-by-step use of raised images became the starting point for an animated component of the site. From a teaching point of view, Caro's method had relied on some all important acts of intervention: the deconstruction of the image into key areas, the emphasis or avoidance of certain details, the flow, over time, through different levels of understanding. Animation could offer a similarly selective, time-based movement through the artwork, which would be supported by text as the touch tours were supported by dialogue. Obviously, animation would be of no value to a blind user. But we were convinced that gentle transitions from one simplified element of an image to another could be an important tool for those with limited useful sight, and significantly, for those sighted people who's own understanding is crucial to the understanding of the blind people they assist.

Any Web site which cropped, dissected and animated the works of Matisse and Picasso was bound to overstep copyright restrictions by some distance, so we created a mock-up version of one work, Matisse's The Moroccans, that could be used to demonstrate our intentions to the artists' estates. To my mind, the result betrays the struggle to reconcile the unique demands of i-Map with the wealth of Web accessibility advice. The page is split into two frames. On the right, a text description of the work contains keywords and phrases, which, as JavaScript links, act as triggers to corresponding animated sequences in the frame on the left. The image section was built in Flash and fundamentally serves to pan and zoom in on relevant areas, mask peripheral detail and occasionally introduce diagrammatic elements to make explicit some of the more subtle structures at work in the painting. As such, the image of the *Moroccans* is virtually ever-present, but other than as a point of departure, is only ever revealed in its entirety at the end. Above the image itself, a control panel of sorts provides the user with sound-coded buttons to view the full image or to return to the main menu. A third symbol, a flashing M, appears at points in the sequence where a raised image is to be used ("Minolta" is a trade name for the raised images).

It is a design that tries to be all things to all users. For example, a blind person, reliant solely on text and raised images, would nonetheless have to contend with the distracting presence of the animation – and this, presumably, after having to download a Flash plug-in for which they had little use. Meanwhile, the user with some sight was expected to follow the text, the animation and the raised drawings simultaneously – a challenge even for a (fully sighted) person. Dan Porter was keen to keep the design uncluttered to give the user the best possible chance of seeing what was onscreen, but this would be impossible for as long as we tried to contain the whole project within the same interface. Perhaps somewhat surprisingly, it had also not occurred to us at this stage that a user that required noises to alert them to the presence of buttons would be highly unlikely to be operating a computer with a mouse. But the steep learning curve was steep, and, as in any design process, identifying these early faults would be prove crucial to the future development of the project. Furthermore, this prototype version of i-Map helped to convince the families of Matisse and Picasso of the integrity of the project, when we visited them in Paris to discuss copyright permission in February 2002.

Having been granted permission to proceed with i-Map, Dan was at a natural point at which to reassess the usability of the site. It also meant , we were now forced to think about i-Map in relation to the Tate Web site as a whole. Beginning with a typical page from the Tate site, Dan attempted to transform it from a webpage accessible to blind and visually impaired users to one that was ideal. He did not have complete freedom to redevelop the page. He had to maintain some design continuity with the rest of the Tate site and the sponsor's logo was also obligatory, but in some ways these were helpful constraints. Having removed the search facility, an animated graphic and the bold strip of colour at the top of the page, a working version of the page was tested by Paul Porter, a blind Web expert at the RNIB. We asked him to list any obvious design faults but interestingly, this information proved difficult to extract. Paul tended to set acceptable rather than optimal usability at the top end of his scale, presumably a mindset of one accustomed to using products designed in the sighted world, and in particular to Web sites 'made accessible' post-design. However, text reader software is a fantastic tool for disclosing inconvenient or badly prioritised content. Unlike the trained eye, it cannot skip automatically to the most relevant content, and after hearing the synthesised voice recite Tate's standard navigation links a number of times, the decision to shift them from the top to the foot of the page was an obvious one. Testing with a screen reader also allowed me to realise, belatedly, that blind users relying only on text and raised images would be best served by pages devoted to that purpose. i-Map's themed pairs needed to be presented in two formats, text only and text and animation, and it was important that at the gateway to these sections we made the intended users explicit.

Meanwhile a pattern of working had emerged. We would begin by storyboarding a journey through the artwork, which would provide us with a sequential blueprint both for the animation and for a set of raised images. Given the range of artistic approaches within our four pairs, it would have been impossible to find a single methodology to produce a narrative for each. Where we thought a sense of the construction of the painting was important, like the Picasso still life or *Nude in an Armchair* 1932, we used a layer-by-layer approach. *The Moroccans* 1916, on the other hand, is a painting in which three very noticeably distinct areas of detail exist; we wanted to examine each of these areas in turn before discussing the spatial conflicts between them and the shapes and colours that gave them unity. So in this case it made more sense to begin with a circuit of the painting, selecting details to act as stepping-stones in a clockwise movement around the work. Our treatment of detail was in itself an important point of discussion. As is it was important that i-Map allowed a visually impaired person to reach the same level of understanding as any sighted person, we could not omit intricate or obscure detail if we deemed it significant to the overall interpretation. In both the animations and the raised images we would find ways of clarifying difficult-to-read areas without damaging integrity of the work, redrawing them in a bolder, simpler fashion, and making sure to state in the text that the artist intended an element of ambiguity. In our storyboarding sessions we were also careful that such attention to detail was not immediately followed by a much wider view, as dramatic leaps in scale would be likely to cause confusion.

Of course, these decisions had a significant impact on Dan's approach to the animation. While his first attempt had, quite literally, revolved around a single, electronic reproduction, it became clear that this could only ever allow us to work on the surface of the painting. For us, the artwork itself only represented the most detailed, final frame or the top layer in a sequence. If, using some creative license where necessary to complete outlines and forms, we could break the image down into a series of independent layers -each wholly legible in its own right- we could build up to a final image through a series of coherent, unobstructed steps. Just as the first raised drawings in a set needed to be intelligible independently of last, so the animation turned the image into a fully articulated object: made up of overlapping elements with the potential to be moved, developed and seen independently of each other. This approach was applied to all the works, but is perhaps best exemplified by Picasso's still life, *Bowl of Fruit, Violin and Bottle*, 1914. Though a painting, it is very much based upon Picasso's use of collage, itself a layer-upon-layer technique. So it seemed very appropriate to give the user the experience of building up the work by arranging cut-out elements on top of each other. The image builds piece by piece as the user hits the spacebar, while the use of sound acts as confirmation that they are affecting change.

Entirely recreating the works had the added advantage of enhancing their clarity. They were built in Flash rather than painted with oil on canvas, and so, for our purposes, they benefit from bolder colour and cleaner lines. There were other important changes from the prototype version. The simpler interface was now fully operable with a keyboard, although it remained possible to use mouse if preferred. And the relationship between the animation and the text had switched from being text-led to animation-led. Key points in the animation were now catalysts for changes in the text, rather than vice versa. Technically, this was a much more reliable method than the first, which, even in the prototype had been prone to faults and cross-platform difficulties.

The preliminary storyboarding involved detailed discussion between us and mutual clarification of the key points we wanted as the focus of each work and how they'd be revealed. So the final storyboard reflected the narrative structure of the text which Caro wrote in response to the completed animation. The text was passed between us for editing and improving, but you'll notice that its language, tone and flow is very different from anything else on the Tate's Web site and from most written interpretation one sees in a gallery. This is because in order for aural information to be absorbed, particularly when it's being spoken by a synthesised voice that lacks cadence and emphasis, sentence structures must be short and simple, information has to be delivered systematically and incrementally, and meaning has to be clear even when delivered in a monotonous deadpan voice. Moreover, using common written terms like 'the latter' or 'the previous point' only confuse because the ear cannot dart back over previous text to remind the listener what the latter point was. This commitment to clarity meant that i-Map's texts are very long, which goes against our instincts as curators and Web designers. However, visually impaired people are used to listening attentively and prefer detailed explanations to hasty summations that often fail to put them fully in the picture. The text only version is even longer because of course each raised drawing had to be introduced before it can be put to use in the service of the overall exploration of a work.

Obviously, in building a project specifically for visually impaired users, a Web designer would be foolish to ignore the now well-established tenets of Web accessibility, and we like to think we could claim to have complied with those. But we would also hope that i-Map illustrates the limited significance of that claim when providing Web access to visual art. In testing the site with blind users particularly, we saw confirmation of the fact that technical accessibility does not mean intellectual accessibility. Indeed, art museums are frequently reprimanded for forgetting to include ALT tags, for example, to identify an image. But in simply naming an image, a museum does little more than give the blind or visually impaired user an idea of what they are missing. In following accessibility guidelines alone, we do not provide access to artworks, but to information about them. And although as sighted people we are often barely aware of it, much of that information is deeply rooted in the visual experience. The Tate Web site, for example, provides access to a vast range of texts and data about artworks and exhibitions, and a blind user should have no problems in navigating it. But the amount information coded in such a way as to include the visually impaired is minimal. It is obviously not feasible for a visual art Web site to author all its from a visually impaired perspective, though its arguably a good exercise for art history students. It is only through specialised provisions like i-Map that visual art on the Web can become truly accessible, so museums and galleries need the time and the resources to make these available.

From GUI to Gallery: A Study of Online Virtual Environments

Stephen Lawrence Guynup, Georgia Institute of Technology, USA

http://www.lcc.gatech.edu/idt/

Abstract

This paper began as an attempt to clarify and classify the development of Web3D environments from 1995 to the present. In that process, important facts came to light. A large proportion of these sites were virtual galleries and museums. Second, these same environments covered a wide array of architectural interpretations and represented some of the most cutting-edge work in the medium. It became clear that there is a relationship between galleries and virtual environments. At a fundamental level, both are information spaces. A primary difference is that the Web3D environments are bound to the computer and currently limited by mouse and screen. Factoring this in, we merge the GUI and the Gallery to create a native foundation for the development of virtual space. This paper discusses the relationship of GUI and gallery, and the impact of mouse and screen, and then showcases the exploration of several on-line virtual museums and galleries.

Keywords: Virtual Environment, Web3D, GUI, Virtual Reality, VRML, Shockwave 3D

A Foundation

Introduction

The digital landscape is littered with failed virtual environments. Questions of their native design remain unresolved. Examining an architectural genre, one similar in functionality to virtual space, would lay the foundation for a deeper understanding of the nature of medium. The question then is "What type of physical space best represents the native form and function found in virtual environments?" In terms of flexibility of design and overall function, one choice clearly stands out – the modern museum. Unlike other spaces like restaurants, bathrooms or garages, the purpose of a museum is not a physical one. It is an educational, often experiential, even spiritual one. With a degree of spatial freedom found nowhere else, museums structure space and house the widest possible array of objects and information.

It is this array of objects, and the ability to access them, that brings us to a second similar design genre –that of the graphical user interface (GUI). GUIs are typically two-dimensional in nature and are usually designed to suit the affordances of mouse and screen. Moving the GUI into three dimensions and making it immersive is the functional imperative of virtual environments. Here again the construct of the museum emerges. A comment by William Rubin, former Director of MOMA, sets the stage: "Museums are essentially compromises ...Their weakness is that they are necessarily homogenized - emptied of all connotations other than art, and that is an artificial situation." To the author, this implies the same processes and the same "necessities" that drive interface design. The museum is the original virtual interface. The outcome of these statements is the belief that the future of virtual design is a merger of GUI and gallery.

Historical Structures and Philosophic Overview

This function of displaying objects and information is the key connection in the historical development and definition of the museum. From the early groupings of precious objects in ancient tombs, temples and churches (Schatzkammen) and the collections of art that hung in the galleries of English country homes and French Castles, to the Italian studiolo, the German Wunderkammer and Kunstkammer, there is an order purposely imposed on the objects and space. It is a space designed to support information, in access and understanding.

In the abstract, a virtual environment can be interpreted as a pure information space. Within it, the functional imperatives of the GUI merge with the appearance of physical space. Like a GUI, information is divided, categorized, placed within a narrative framework, and then presented to the user. The user is then given paths to follow to access the information. This mirrors real world galleries. Space itself facilitates the access of art. At their best, museums and galleries are flexible spaces that uplift, amuse, educate, classify and present information with a degree of spatial freedom found in no other structures. In this sense, the purpose of the gallery is the same as the GUI. For the designer, art = information in its widest array of configurations.

This is somewhat different from the earlier connection made between virtual environments and museums. "Museums are ideal candidates for hybrid cyberspaces" (1990 Kellogg, Carroll, Richards), and their proposed "Natural History Cyberspace Museum" is a fully immersive installation complete with goggles and gloves. Through it, visitors can pretend to be dinosaurs such an Archaeopteryx. While this application of virtual space is clearly valuable and could be very useful to educate the public, it does little to express the affordances of virtual space. Rather than apply virtual space to the museum, we apply the museum to virtual space.

Secondly, and perhaps controversially, we equate Web3D environments with virtual environments. We do this not as Ken Hillis does in "*Digital Sensations*", completely removing technology on philosophical grounds. Instead, we see the future as undefined. Therefore, the legacy of currently widely available technologies is liable to hold a greater sway over this future than imaginary plot-devices such as the holodeck. Within current Web-based environments in mind, we accept the limitations of mouse and screen.

Interpretations

A diverse series of interpretations of 3D space have been produced. Most attempt to copy real world spaces exactly. In them, we see faux track lighting, cushioned benches and display cases. Real approaches often go back in time, like the "*Natural History Cyberspace Museum*." Opposing them are those whose work represents a rebellion against the limitations of real world architecture. The "*Virtual Guggenheim*" by Asymtote is such a space. The structure was envisioned to morph between three different forms, while all the while rooms flow like blood vessels through it. Still others, like the early works of Stan George and Daniel Wise, resemble abstract data spaces envisioned by sci-fi authors. Often visually stunning, none of these approaches has yielded the results predicted by their creators. Each suffers from the imposition of philosophies that do not take into account the affordances of this new media.

Hillis (1999) points to "the promise and hype of VR and ITs more generally is part of an ideology of the future, produced in an amnesia and loss of history that forgets the broken promises of past technologies such as the 'universal educator' [TV] and 'too cheap to meter' [nuclear power]. Whereas Hillis dismisses the virtual entirely, we take a practical approach and look to design in a useful medium, based on facts and not what Hillis deems "*wishful thinking.*" The goal of this paper is to briefly discuss these failings of virtual space and then provide examples that point toward a native virtual architecture based on the GUI and the gallery.

The Impact Of Mouse And Screen

Examining Failure

The reasons for failure are diverse. The literal copying of the physical space into virtual space has been less than successful. As we address these failures, we reinforce the previous assertions of gallery and the GUI. We can begin with a common misconception, that *virtual reality works just like physical reality.* This idea completely fails to account for the hardware through which we interact. Screen-based, mouse-driven, virtual environments (SMVEs) utilize a vastly different interactive methodology. Two issues that quickly become apparent are navigation and interaction.

Unseen Feet

Navigation issues arise on the ends of the spectrum; wayfinding in larger environments and short, goal-oriented, precision movements. To understand precision, one can look to video game design. Games

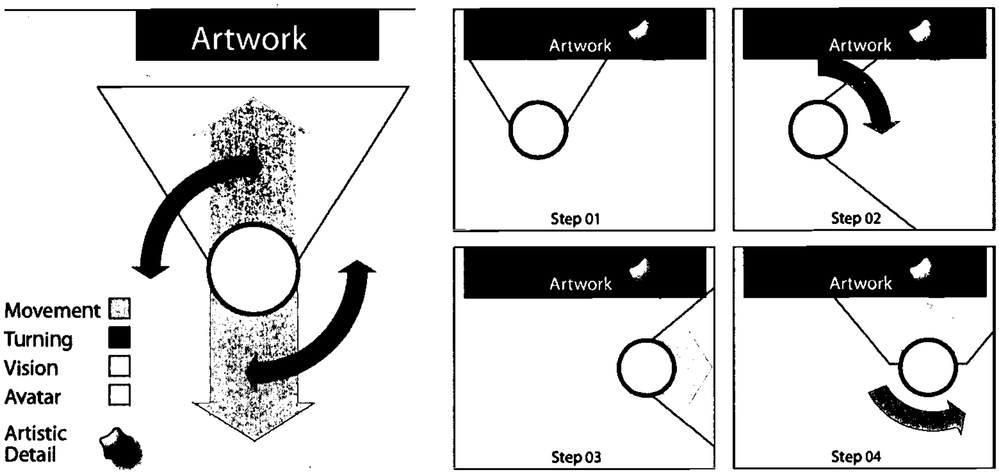

Figure 1: S. Guynup, Diagram of Short, Goal-Oriented Movement

that require precise positioning in their environment use a third-person viewpoint. Less immersive than the first-person viewpoint, the third-person perspective allows a uniquely powerful navigational advantage – you can see your feet. As simple as that sounds, it makes all the difference. Like a cursor in a GUI, feet show the absolute position of the user in the environment.

More precision issues lie in the standard movements allowed by most Web3d applications: move forward/backward, turn left/right. Compared to the flexibilities of human mobility, the two-button mouse is quite handicapped. The issue is evident when the user zooms in to see a work close-up, and then wants to see another section or detail of that work close-up. The diagram below shows the unwieldy process by which most users try to see the detail. Note Step 03: users need to position themselves in front of the detail without seeing it.

Again, video games have addressed this issue by allowing the user to strafe. Strafing is the ability to move perpendicular to your viewpoint. This side-to-side motion is very useful for first-person shooter games where the purpose is to move quickly and kill. Some Web3D technologies allow for sliding, yet even the most experienced users fail to change their navigation preferences for such seemingly short actions.

A last, hidden issue relates to the purpose of the space. The design and action in a game environment are driven actively by the computer and the narrative. In the (usually) more subdued genre of museums, the addition of strafing navigation ability is helpful, yet not as effective. In the museum, Art drives the user's action. The mindset of users examining art makes them more sensitive to the limitations of the interactivity.

Wayfinding, the cognitive process of navigating and interpreting the environment, carries a subtle problem. Typically, it takes time to understand the spatial mapping of any environment, real or virtual. We seem to forget how easily we get lost in new surroundings. Factor in the lack of nuanced cues like physical wear, and the problem grows. Wayfinding becomes even harder as designers, free of construction restraints, build large spaces. Free of natural physical laws, they build using their own site-specific logic. (For a secondary reference, look how often we get lost in a two dimensional Web site.) Over time, the user does learn how to interpret the layout of the space. Unfortunately, unlimited time is something we do not have.

Mouse vs. Body

Interacting with virtual environments carries issues similar to navigation in Web environments. The

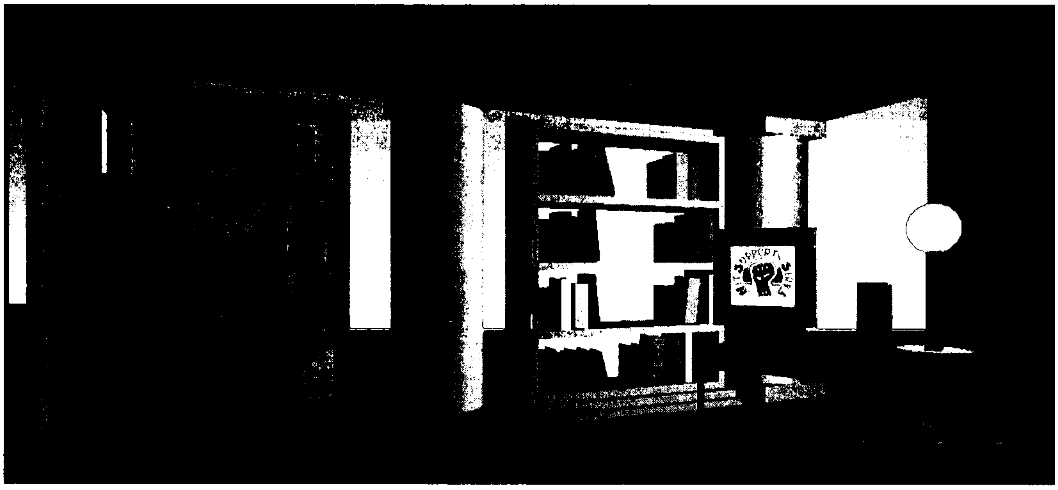

Figure 2: J. Sonstein. Library in "Town Square at Ariadne" http://ariadne.iz.net/~jeffs/

mouse is not the hand. This sets up the following dynamic; Visually, we cue to items that we interact with in the real world. If the interaction is beyond the scope of the mouse, we can frustrate the user. Typically, the question they ask is "why bother?" Below is a unique example.

Each book links to a topic in the Yahoo® database; the computer is a link to a Bosnian info site. Metaphorically, the layering of informational structures (i.e. books and computers) inside a virtual space is thought provoking. In terms of functionality, it is less than optimal. Books and computers were created for use by human hands. The mouse was developed for two-dimensional interfaces.

Leaving the Real

The reasons for merging GUI and gallery are becoming apparent as mirrored reality approaches falter. As technology advances and allows for more natural manipulation and navigation, we realize that the greatest empowerment may still lie beyond the current notion of reality. Take, for example, the fundamental mechanism of interactive design – the button. Does the button exist in nature? Do birds and bees press buttons? This abstract means of control constructed with finger operation in mind and placed within arm's reach has become to us second-nature. Recognized by repetitive cause and effect, the button becomes intuitive. Even when the mechanism below the button is beyond the user's comprehension, such as physics and the codes that make my computer keyboard operate, the action feels natural. As we move toward a native virtual design, vast possibilities emerge, possibilities that stem from a legacy of mouse and screen. The next section moves into abstract territory, presents issues inherent when reality is left behind, and then offers harbingers of future design.

Harbingers of Native Design

A Realm of Pure Information

Authors like William Gibson and architects such as Michael Benedikt left reality behind. Space as interface, as database, a "realm of pure information", this was the desire of early virtual reality theorists. Abstract visionary environments were designed for the human mind, not the body. Information in graphs and charts surround the user. The philosophic papers of this genre are deeply moving - these environments, these dynamic hyperspaces are beautiful works of art. Yet, overloaded with imagery and force fitted with data, their functionality is forgotten. Despite the ability of three-dimensional space to hold vast amounts of data, it is important to remember that; "Less is more".

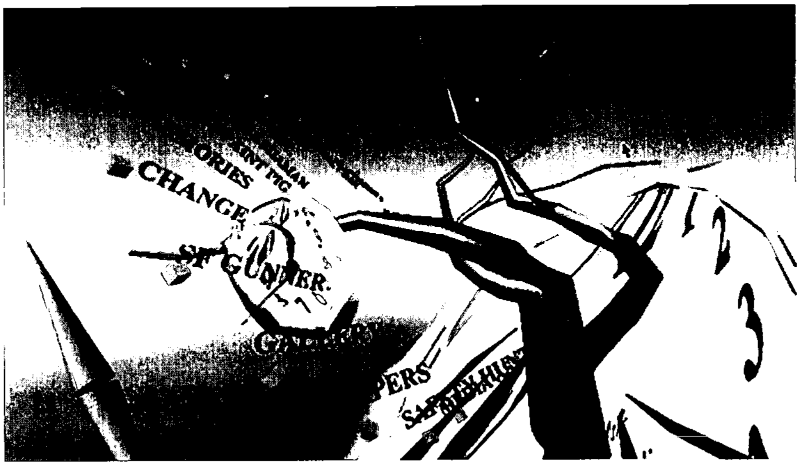

Figure 3: S.Guynup. Surreal Menu

Data representation in three-dimensions is difficult because spatial-visual information generated by the space does not support and often contradicts the data the developer wants represented. Perspective and shading throw a wench into our ability to compare and quantify. Furthermore, the informational narrative created by moving through or rotating the space offers little practical benefit for data that is not inherently spatial or placed with that in mind. To delete the spatial information is to return to two-dimensional space. This then becomes a dilemma of purposes: navigatible space requires visual-spatial information that cues the users, allowing them to move through it, but data representation seeks clarity and comparability.

We can summarize the failings of prior works in two general categories: realistic environments that did not compensate for the affordances of mouse and screen, and data driven environments which could not structure the information to the required level of clarity. The former is a function of technology, the latter of purpose.

To find solutions is a two-fold process. In the next section, *Towards a New Interaction and Architecture*, we'll briefly address native navigation and interaction techniques. For the second part, the question of purpose, we move from realistic or data driven ideals and return to William Rubin's issue – that "museums are ... emptied of all connotations other than art." We explore connotations and connections in the section entitled *Examples of Exploration.*

Towards a New Interaction and Architecture

To address the root issues of user interactions, we must design objects to be user friendly. Size is often the simplest fix: larger objects are easier to see, easier to click on and proportionately easier to correctly navigate to. Smart objects are possible, and they occupy a range of possibilities. Most Web3D languages allow for "Billboard" objects. These flat panels or even full 3D objects automatically turn to face the viewer. If the panel contains visual or interactive elements, this is an easy way to allow the user to read or touch it head on. Beyond turning on the y-axis, the objects or panels can gently slide into the path of the user. A more aggressive approach is having the interactive object grab the user and position him or her as needed. The effect of this is somewhat jarring to the user and, if used, ought be supported by the narrative of the space (i.e. the designer can provide an alternate reason for moving the user).

A designer can also break from the interior of the space and place information into a heads-up display or even into a side HTML frame. This pop-up can be based on user positioning or be activated by a button on the user's "dashboard". While a level of immersiveness is lost, the ease by which the user can interact is greatly enhanced.

When designers remove items like chairs, tables, lamps from a screen-based, mouse-driven virtual environment (SMVE), a sense of emptiness appears. The effect is emotional as the user has few pre-existing cues to guide the relationship with the space. Real world items also form a basis for comprehending the size and layout of the space. In appropriate amounts, these elements cue the user's navigational behavior (in larger - or even just realistic amounts they can be cumbersome). To compensate, some virtual architects take a deconstructive approach. They add visual interest and navigational cues by breaking the flat planes of wall and floor. Slight angles and wire frames can reinforce the user's sense of three-dimensional space. Some authors add background animation to give "life" to a space.

Patrick Keller's Alternate Fabrique shows the classic architectural interpretation of virtual structures. To generate a sense of space that appears full yet lacks benches and lights of the realistic spaces, Keller breaks and bends his galleries, extends faux structural elements outwards. Lines stretch like taut ropes and help define forms that are weighted yet weightless.

Often, designers use bold colors and oversized visual elements to fully fill their virtual space. This effect, usually fun and visually engaging, plays off the user's previous knowledge of two-dimensional animated cartoons. There is a second benefit to scaling elements in the environment upwards. Large elements are typically easier to see and interact with. Large elements give the user a larger area for proper positioning and can be a solution to problems of precision navigation

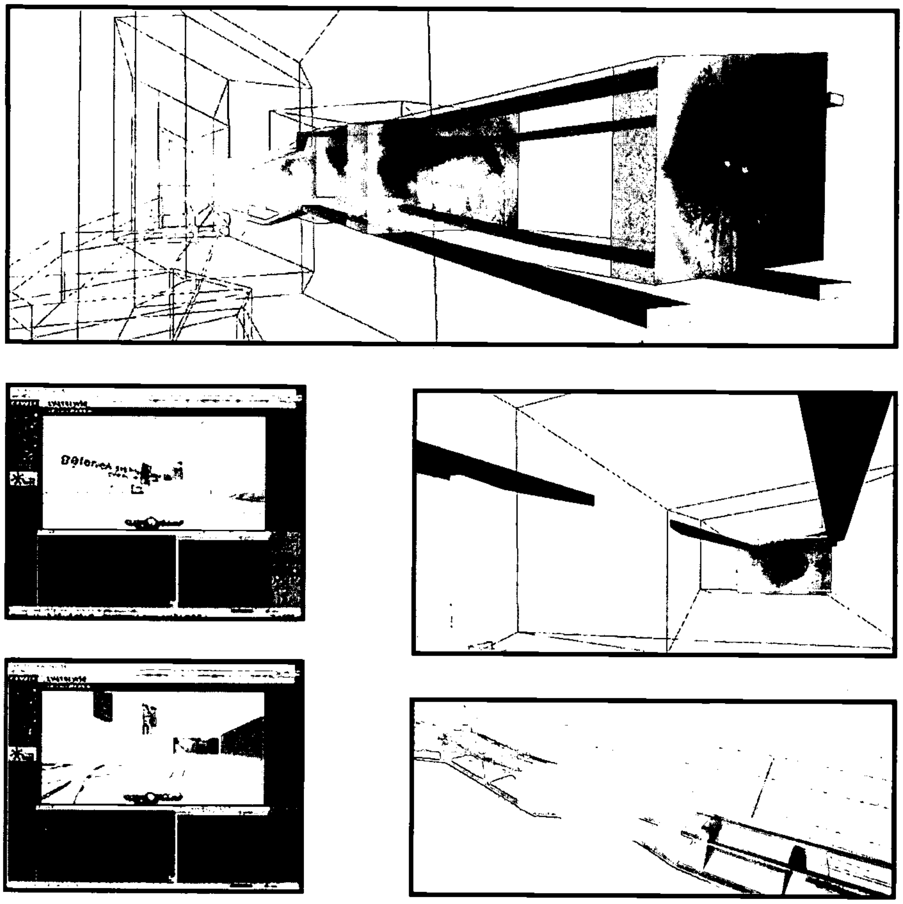

Figure 4: P. Keller. The Alternate Fabric gallery "Expo 1:prothese digital"
http://www.fabric.ch/La_Fabrique00/prothese.html

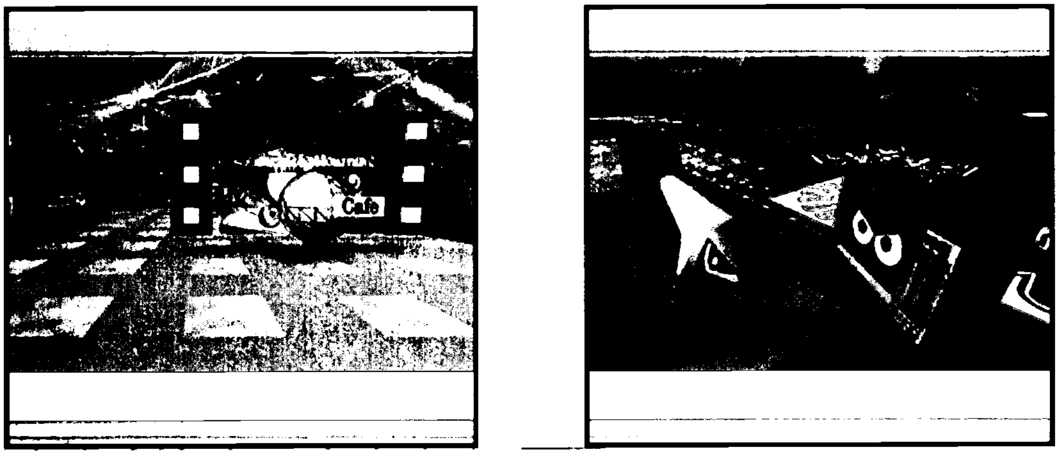

Figure 5: an information space from VGTV's Global Learning Prototype ,
K. Dudesek, founder. http://www.vgtv.com/p,500029,1.html

Both of these approaches represent a lateral expansion of reality. To understand this, compare the real world environment of New York City to that of the Arctic tundra. Visually, they are extremely different, yet the same laws of nature and physics apply. Gravity is always active, and the ground is solid. In virtual environments, these rules are flexible; they are at the whim of the world builder. At a deeply fundamental level, there are abstract concepts governing reality that are subject to review. In real environments, objects are persistent and shared. In the virtual, these rules are by default, reversed. Objects are typically unshared and non-persistent. While this is seen as commonplace within a GUI, the long-term impact on the design of virtual space is unknown.

As we exploit the native characteristics of SMVEs, a path becomes clear. We merge the speed and clarity of two-dimensional interfaces with the immersive, experiential aspects of three. In structuring space as subsets of explorable data, we can group and display information in new ways. We can add a level of connection and to some extent the connotations that were previously unreachable.

Examples of Exploration

The following museum built by Christiano Bianci typifies the subtle border of lateral expression. It has a strong and controlled color scheme and presents its exhibitions within wire frame boxes. At the same time, real world new media trade shows and art exhibits have a very similar sense of place.

Stairs could have been used to connect the three floors of this virtual gallery, but in these environ-

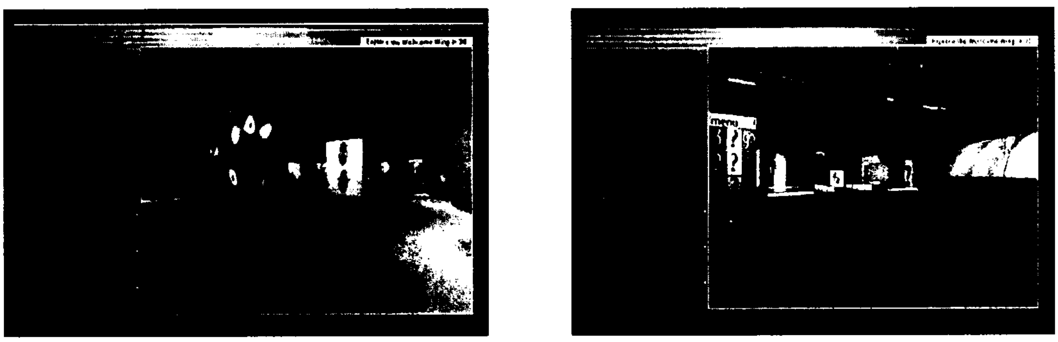

Figure 6: C. Bianchi. Interior of "The Science Museum"

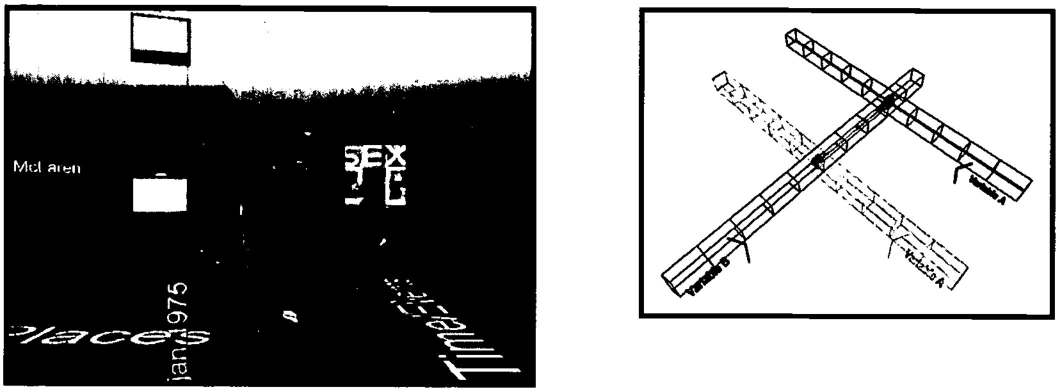

Figure 7: C. Bianci. A database as gallery structure

ments, stairs are difficult to use. So instead, we teleport, using what appears to be a kiosk (left image). Note that while it resembles a kiosk, it utilizes a GUI menu structure that in fact even says "menu".

The Navigatible Database

This discussion begins, surprisingly enough, with a space that responds to the user's movement. Although not typically seen as mobile, databases possess a formal order that is navigated and manipulated. We can transcribe the database into a three dimensional form. The example below by Christiano Bianci does just that. As the users move through the space, the data appears to them. The left image is a screen capturel; the right is an illustration of action.

Image this structure within a database table. Related rows and columns slide to present the intersection between the data types. The intersection is given in full detail and the corresponding row and column is given in outline. To navigate this environment, we enter it and we exist at the intersection. As we move, so does the row and column. Data in full is presented around us. The upper right diagram illustrates the layout. Variable B is a timeline; variable A are the members of the band known as the Sex Pistols®. The choice was made because of the diversity and availability of information about the band. Within each database unit, we find text, images, sounds and movies.

Relating space to information can be done on many levels. The gallery on the following page does just that. In the distance, it is an object, a spinnable timeline of the 1950s to the 1970s. The goal was to create a relationship between the distortion of perspective caused by the user's vantage point and the data. Above, it highlights the differing viewpoints on the 1960s. In concept it is as simple as saying, "In the eyes of 1950s, we see the 1960s as 'X' while those looking back from the 1970s see the 1960s as 'Y'."

Secondly, the timeline object can become an environment. The user can move into the space and navigate towards the center, the 1960s. In doing so, a narrative is produced that colors the user's opinions of that controversial decade. Depending on which entry point, either from the 1950s or the 1970s, a distinct outlook is formed. In this sense, we mirror the real world: our previous experiences deeply influence our opinion of what follows. Lastly, to the left is a heads-up display (HUD) that tracks the user's position in the timeline. Visually, the user is able to relate his environmental position to the abstract point in time.

Divisions of Information: Kiosks within a Menu

Information in a virtual scene can be brought to the user in any quantity desired. This ranges from a single audio file or image to entire worlds that hold a variety of informational objects. In the middle range is an installation, or to use a more functional analogy, a kiosk. A virtual kiosk is a complete unit of information accessible to the larger whole via a menu.

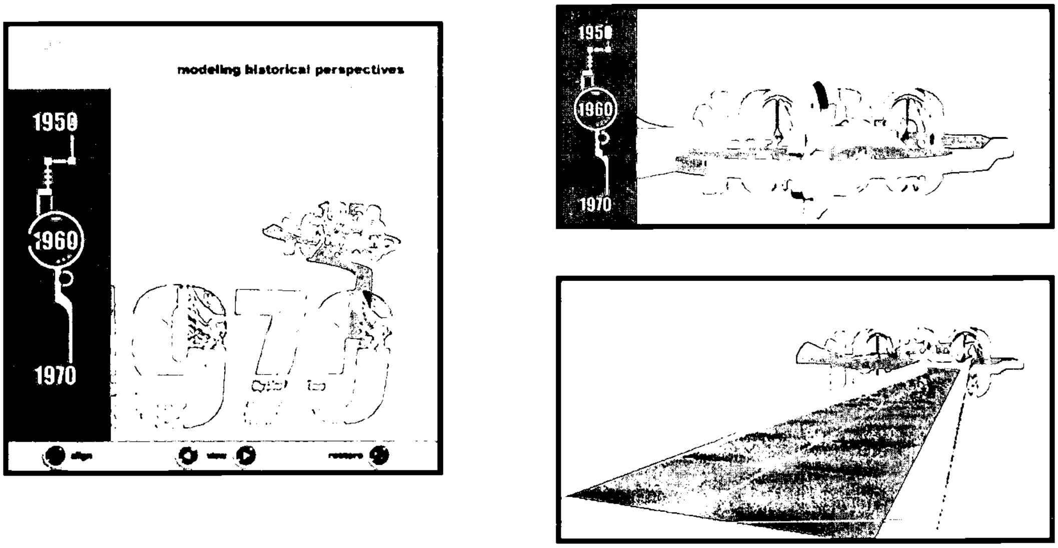

Figure 8: S. Guynup. "Timeline" an immersive pathway

Clicking on the red buttons in the menu arm gives the user access to the other installations / kiosks. Five separate, yet related informational structures are within easy reach of the virtual visitor. With a single click, users can see the information they came for and decide if they wish to explore it in detail. This could have been created in a single large space holding all five installations. Does the users' travel between installations have value, or is it (in this instance) time consuming and possibly disorientationg? The users may get lost. If they demand quick access, as most do, they will grow frustrated with the one-space design's inability to meet their needs.

And of a New Art

We have addressed issues of usability and have hopefully provided practical information on the nature of virtual design. Looking forward, we see not only new structures and interfaces, but new art as well: an art that will allow us to explore the relation-

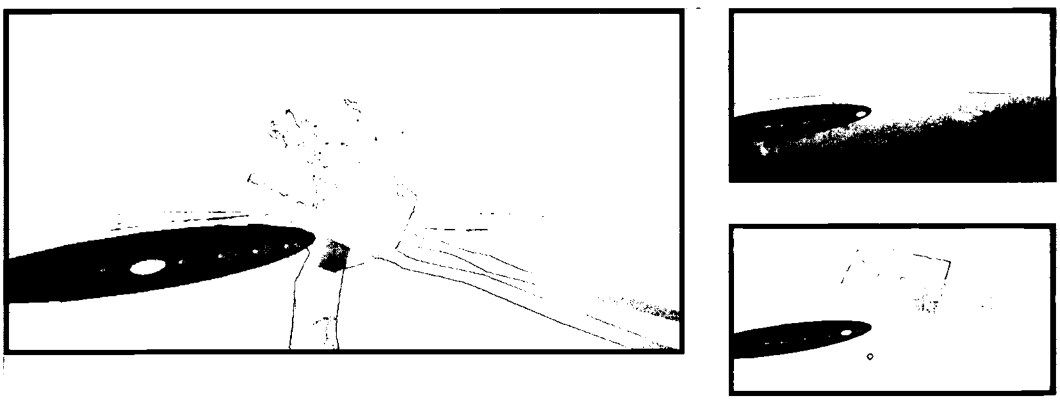

Figure 9: S. Guynup. "Untitled Memories" virtual installation

Figure 10: M. Clifford, "The Aleph" immersive paintings. http://www.rre.net/welcome/matrix/public/

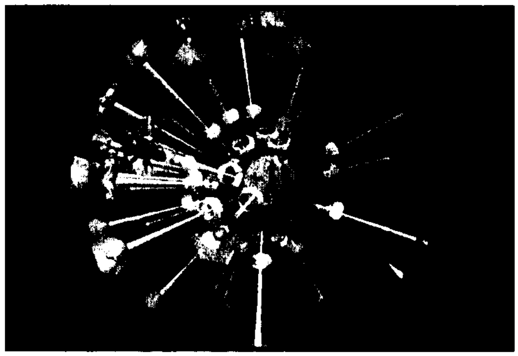

Figure 11: J. Klima, "Glasbead" dynamic interface / musical instrument. http://www.cityarts.com/glasbeadweb/

ships between our bodies, minds and world. The variety of works and the galleries that support them is vast. Much has already been done.

The painter Maurice Clifford has translated his abstract art into animated, almost hypnotic flowing sculptures. In John Klima's Glasbead, art and interface design merge into uncharted territory. Spun by the user, the hammers freely rotate around the center. When a hammer collides with another, a tone is made, and so the physical rhythms of spun hammers create a new musical form.

Earlier we discussed the gallery structure created by Patrick Keller. Conceptually he had designed it to hold virtual works of art. The virtual six artists invited to show in the space were confronted with a unique problem. For artists accustomed to infinite space, the real-world-sized rooms were in effect very small. The smallness of space became the starting point for many of the artists. Taking a cue from the cramped quarters, Andy Best filled his small space with a large fiery animated animal. Steve Guynup added a single button. When pressed, it caused all but the entrance and exit portals to flitter away. With walls removed, the users found themselves in an abstract Chinese garden. In addition, as this was a multi-user space, the users could now interact with others in wholly separate visual environments.

The most ingenious tack was taken by Christiano Bianchi. His work consisted of boxes that float in space at a height that matched users' heads. With collision turned off, the users can actually enter the box. Upon entering a box, users found Christiano had altered their walking speed, making them move more slowly. As the primary tool for judging the size of abstract space is the speed through which one moves, he had in essence created a larger space inside the box. He programmed it so that infinite boxes were placed inside each other and in each, the users moves a little more slowly. Adding letters to the boxes, he entitled the space, "Infinite Babel"

A Goal

Towards a Deeper Understanding

The goal of the museum is "not to create subject matter mastery, but to accomplish certain appreciative goals—to awaken interest, to broaden perspectives, to induce deeper understanding, to enrich aesthetic sensitivities, and so forth" (Lee 1968). To do this, we can look beyond the real and the dreamscape to the new.

The gallery below supports two installations. The center holds a sculptured body circled by NATO and Serbian emblems. In the middle of the figure is a map of Kosovo. The interaction is simple. The art attacks the user. NATO and Serbian forces fire lasers and the user can fire back. The only damage

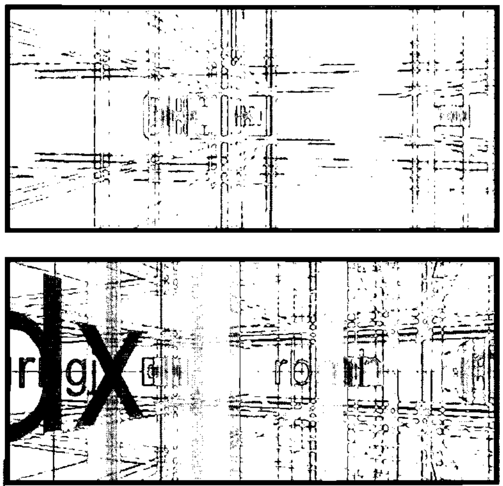

C. Bianci. "Infinite Babel"

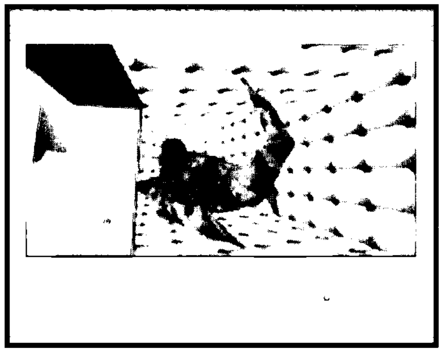

A. Best. "Animal"

Below: S. Guynup. "Digital Zen Garden"

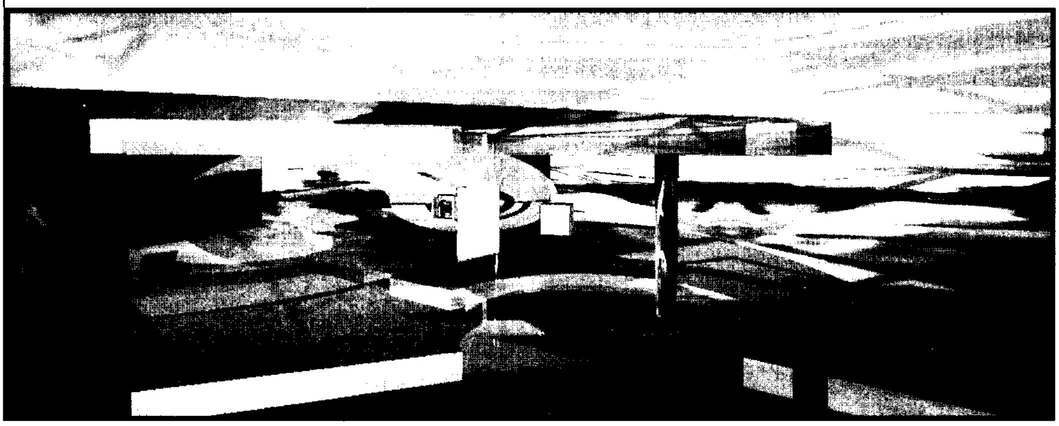

Figures 12-14: Examples of the show at "Expo 1:prothese digital"
http://www.fabric.ch/La_Fabrique00/galerie.html

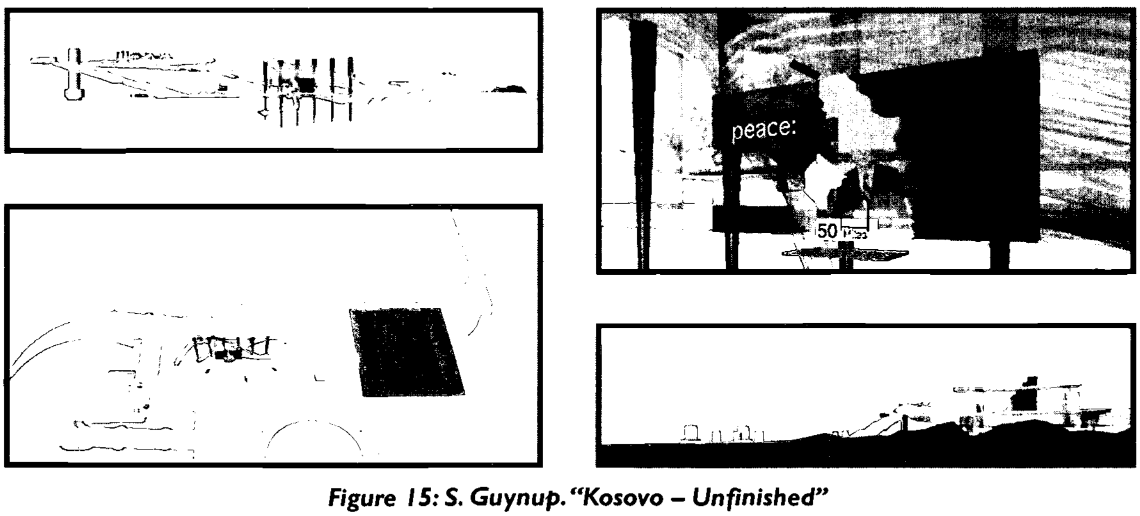

Figure 15: S. Guynup. "Kosovo – Unfinished"

they do is to Kosovo. In the upper gallery, images of the destruction are circled around a piece of rotating text that declares "*...this is not a game.*" As cultures become reared on video games and see war as just another image on TV, it becomes crucial that we use our technology not as trip in time or as an escape from reality. It the end, virtual reality must reinforce the real.

Conclusion

In describing virtual space, we have encouraged a broader perspective in terms of structural design and the end use. As time progresses, virtual spaces on the Web will become more commonplace. To start this process, they must accept the current affordances and limitations of mouse and screen. In creating a new interactive methodology, we also redefine the means by which we can interpret our world and ourselves. Broad new vistas of human experience are not far off, and yet, they are distant enough to be unknown. This paper offers at best only harbingers of that future. All that can be said for certain is that the path to this end lies in the merger of the GUI and the gallery.

References

Asymtote. "Virtual Guggenheim" – Asymtote http://www.guggenheim.org/exhibitions/virtual/index.html

Benedikt, M. (1991), Cyberspace: Some Proposals. In M. Benedikt (Ed.) *Cyberspace: First Steps*. Cambridge: MIT Press, 119-224.

Benedikt, M. (1991), Introduction. In M. Benedikt (Ed.) *Cyberspace: First Steps*. Cambridge: MIT Press, 1-25

Bianchi , Christiano. Infinite Babel. How small can you get to find a meaning, 2000, http://www.cristianobianchi.com/html/canal.html

Bianchi, Christiano Bianchi. *"Infinite Babel"*, *"Database Gallery"*, *" Science Museum"* – http://www.keepthinking.it

Bianchi, Christiano. "Beyond the Desktop." An experiment in user interface and responsive spaces, 1999, http://www.cristianobianchi.com/html/beyond.html

Clifford, Maurice. "The Aleph", http://www.benjaminremembered.com/welcome/matrix/public/

Hillis, K.,(1999) *Digital Sensations*, Minneapolis: University of Minnesota Press.

Keller, 2000. Patrick Keller. *"Expo 1:prothese digital* " http://www.fabric.ch/La_Fabrique00/galerie.html

Kellogg, W., Carroll, J., & Richards, J, (1991), Making Reality a Cyberspace. In M. Benedikt (Ed.) *Cyberspace: First Steps*. Cambridge: MIT Press, 1-25

Klima, 2000. John Klima. "Glasbead", http://www.cityarts.com/lmno/

The Meetfactory Andy Best, Creative Director. http://www.meetfactory.com

Mullet, K. & Sano, D. (1995) *Designing Visual Interfaces*, Mountain View: Sun Microsystems

Newhouse, V. (1998) *Towards a New Museum*, New York: The Monacelli Press.

Shaw 2001. Chris Shaw.. Data Visualizations CS 7497 Virtual Environments, Georgia Tech Class 2001 http://www.cc.gatech.edu/classes/AY2002/cs7497_spring/

Sonstein, Jeff. "The Town Square", http://ariadne.iz.net/~jeffs/

VGTV – Karel Dudesek www.web3dart.org

Interfacing the Digital

Steve Dietz, Walker Art Center, USA

http://walkerart.org

Abstract

"Interfacing the Digital" presents work at the Walker Art Center focusing on new physical interfaces, particularly for the presentation of digital art. Some examples include a freestanding revolving door portal for the exhibition *Art Entertainment Network*; a telematic table resulting from an international design competition; and a "temporary autonomous sarai" developed collaboratively by the new media artists Raqs Media Collective (New Delhi) and the architectural practice Atelier Bow-Wow (Tokyo). These and other projects are prototypes for new, interactive social spaces and functions being developed for the Walker's new building expansion, designed by the architects Herzog & de Meuron.

Keywords: digital art, art interfaces, digital curation, Raqs Media Collective (New Delhi), Atelier Bow-Wow (Tokyo), telematic table, Walker Art Center

Background

In the last 5 years, especially following Documenta X (1997), the Whitney Biennial of 2000, and *Net_Condition* at ZKM (2000), there has been at times heated debate not only about how best to present digital and specifically networked art in an institutional context but also whether to do so at all.[1] Not all of the discussion revolves around issues of physical interfaces to such works, but their onsite presentation is a critical concern for both museums and artists—and their audiences

This paper is informed by these discussions, mostly online in the archives of nettime, rhizome, thingist, and CRUMB, but focuses on personal experiences in curating 10 exhibitions over the past 5 years that have included network-based art, including *Beyond Interface: net art and Art on the Net* at Museums and the Web in 1998.[2]

Finally, while working with art that is "born digital" is a special case for most museums, I would argue that many of the issues and lessons are transferable to the digital contextualization of any work in a museum's collection.

The Challenge of Context

One of the challenges of presenting digital art, is that the context and the work are generally displayed via the same means: the screen. How to differentiate between the metadata and the experience? One strategy is simply to open the project in a new window. This was the strategy of *Beyond Interface.*

The complaint from artists about such a strategy is that it creates a curatorial gateway that viewers must pass through before getting to the heart of the matter, the actual experience. If you think of the example of video games, for instance, there may be a narrative introduction to the game – but often there isn't and generally you can skip through it – and then you're in game play mode. If you need/want help you specifically open the FAQ or Help screens, but they are not the main way of starting the game. Even with a painting exhibition, while there is reams of research about the best length, tone, style, etc. for didactics, the working assumption is that most people look at a painting first and then read the label – the help file, so to speak – if they want more information.

Even when net art exhibitions present the artwork first, it is often because the only curatorial context is a list of links, an equally unbalanced approach.

With *Art Entertainment Network*, I tried to finesse this issue by making the interface part of the experience itself – it is a portal to art projects, but to find out more information, you must go to the context (http://aen.walkerart.org).

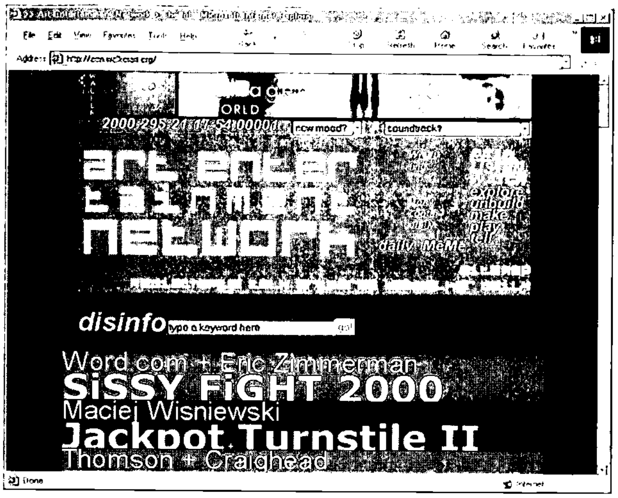

Figure 1: Screenshot Art Entertainment Network

And more recently, with *Translocations*, I integrated functions of various projects into the interface, so that from the contextual pages, you could directly "modify" the text (via an artist project called *OPUS*), send an email about it in any language anywhere (via an artist project called *Translation Map*), create your own soundtrack to browse to (via an artist project called *Translocal Mixer*) or pop up the video window of the *Translocal Channel*.[3]

The Challenge of Ghettoization

Related to the issue of context is ghettoization. There is a conundrum. On the one hand, it can be valuable to provide a focus on a particular set of practices, whether they are photography or performance or digital art. It is easier in such focused contexts to meaningfully differentiate between, say, documentary, fashion, abstract, and conceptual photography, each of which has its own distinct – but intermingled – histories, methods, presentational contexts, etc. At the same time it doesn't make sense to completely divorce photography from the visual arts; to not include it in a thematic show, whether about modernist art in America or America in the modern age.

At the Walker, my answer, as is my wont, is both/and, not either/or. Gallery 9 is a virtual gallery for network-based art. At the same time, with a project like *Shock of the View*, we specifically chose to compare a physical artwork with a digital artwork. In the example below, for instance, the guest curator chose to compare a John Frederick Kensett painting of Mount Washington with a ski resort webcam of the same view, arguing that Kensett was the board-of-tourism-promotional-guy of his day. In another case, I compared Ken Goldberg's *Memento Mori* with the ephemeral *Sisyphus* of Luciano Fabro (http://www.walkerart.org/salons/shockoftheview/).

Comparisons are odious, and in the end, the interesting point is not that digital artist x is as good as artist y, but to bring their contexts into collision and see what happens. Sometimes this is best done in a digital only show. In the same way that Douglas Fogle at the Walker might do a painting-only show, such as *Painting at the Edge of the World*, which explores diverse ideas about painting at the moment, in the Walker's expansion there will be a "mediatheque" devoted to the presentation of new media art (http://www.walkerart.org/programs/vaexhibpainting.html). We will continue, however, to do crossover shows, such as *How Latitudes Become Forms*, the current Walker exhibition curated by Philippe Vergne, which incorporates performance, new media, and film directly into the gallery exhibition (http://latitudes.walkerart.org).

The Challenge of Medium

In this post-medium, post-studio world, the idea of a medium may seem slightly antiquated and naïve. It's all just art, right?

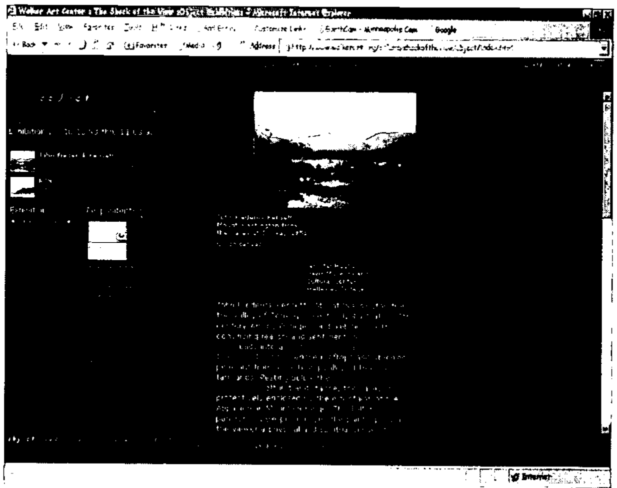

Figure 2: Screenshot Shock of the View

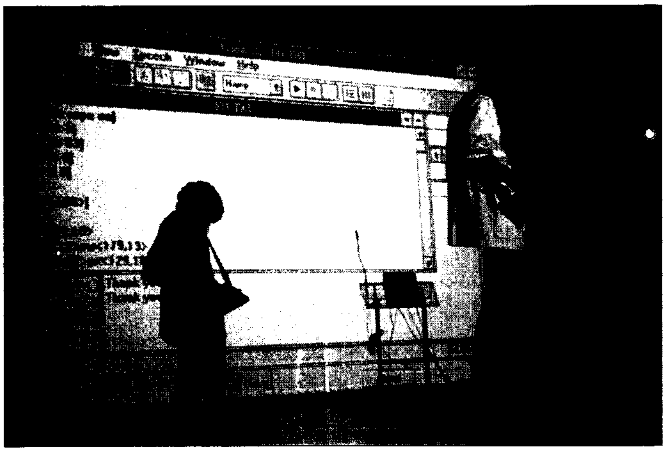

Figure 3: Alexei Shulgin performing as part of the Open Source Lounge at Medi@terra 2000, Athens

Personally, I believe there is a cinema and video practice tradition that cannot be fully subsumed in installation art. The same with photography. And, I would argue, digital art.

Perhaps more to the point, however, is that we tend to talk about most digital art through the prism of visual arts, perhaps with a nod to video art, but in many ways the more fruitful comparisons are with the performing arts. Digital art is time-based, often performative, often ephemeral, often done in/by groups, process-oriented, and so on. The burning issues of collectability and ownership and authenticity take on a whole different tone when viewed against the history of music, its notation system for replaying a core experience that is nevertheless different every time; and the by-now acceptance of live and recorded performances as different but not merely derivative.

With *386 DX*, Alexei Shulgin – shown here in Athens in a show I co-curated with Jenny Marketou for Medi@terra 2000 - plays computer-generated covers of hits such as *Purple Rain*, an inspiration, I would argue, to a whole generation of digital performance artists that transgress the boundaries of the art world's disciplines (http://www.mediaterra.org/mediaterra2000/en.cgi?opensource and http://www.easylife.org/386dx).

The point in terms of interfacing the digital is to pick our models appropriately. Many of the issues of displaying digital work may be better solved working from a tradition of theater (think of object theater in history museums, for instance) and performance than rigid adherence to the traditional gallery experience.

In the Walker's new expansion, we have built in this idea by putting the mediatheque galleries literally in the balcony level of the new performing arts studio. While there is no specific plan to always integrate the spaces, the assumption is that at some point artists will insist on their integration, building out the potential of the conjunction experimentally.

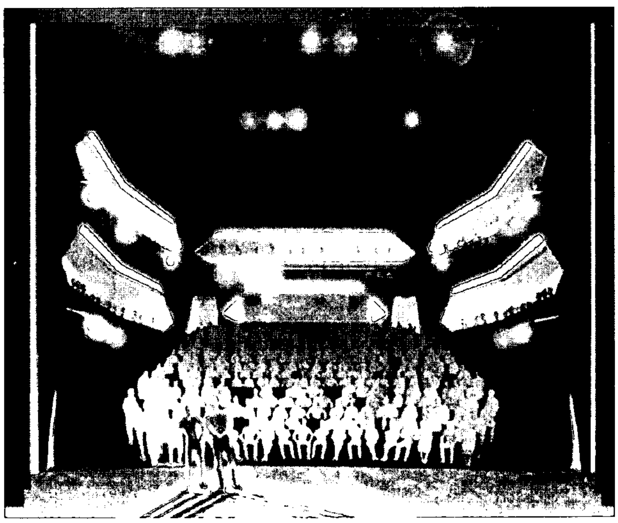

Figure 4: Computer model of the new performing arts studio as part of the Wlaker Art Center expansion. At the back of the 2nd level balconies will be a series of mediatheque spaces.

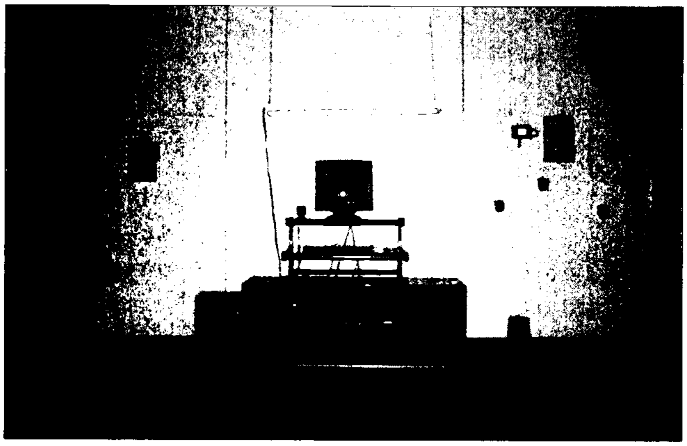

Figure 5: "BangBang," 2000, Bureau of Inverse Technology. Installation view from "Telematic Connections: The Virtual Embrace" at the Oklahoma City Museum of Art

The Challenge of Expectations

One of the reasons the computer and the network have become not only subjects but also means of making artwork is that they are so ubiquitous in our daily lives. That very ubiquity, while it may provide a certain familiarity, also creates a whole set of expectations, starting with "user friendliness."

Try telling Matthew Barney he should be more user friendly. Or Jasper Johns that his references are too obscure.

In this picture of *BangBang*, a work by the Bureau of Inverse Technology, which was presented in a traveling show I curated *Telematic Connections*, the Bureau disturbs our expectations of interactivity. The viewer can do nothing to make the work "happen;" it is dependent on environmental triggers outside of the gallery space (http://telematic.walkerart.org/ and http://telematic.walkerart.org/telereal/bit_index.html). This was extremely frustrating for many visitors, but too often the response was based on a false notion of "good interactivity," not, really, whether they liked or understood the work on its own terms.

Interactivity and user-friendliness are just a couple of the expectations with which we view digital art. On the museum side of the equation, there are a whole set of parallel issues to do with touch/don't touch, how much time you spend with a time-based work, and other learned gallery behaviors. All of these are concerns that both the artist and the presenter must take into consideration when interfacing the digital, but the answer is not always obvious – to make things "easier" to do, for instance.

The Challenge of Infrastructure

There is not really too much to say about infrastructure except that it is absolutely necessary, vital, critical. Yet, in every single installation I have done over the past 5 years, no matter where, the level of support has been less than for a comparable contemporary art exhibition at that institution. Not out of malice or design but largely because of the history of the space and the personnel. Nevertheless, it should be just as easy to plug in to a network connection as it is to electricity, as just one small example. But it's not, and until it is, institutions will only be compounding their problems – and their audiences – when interfacing the digital.

The problem is totally solvable, but perhaps primarily in the design of new facilities. No facility should

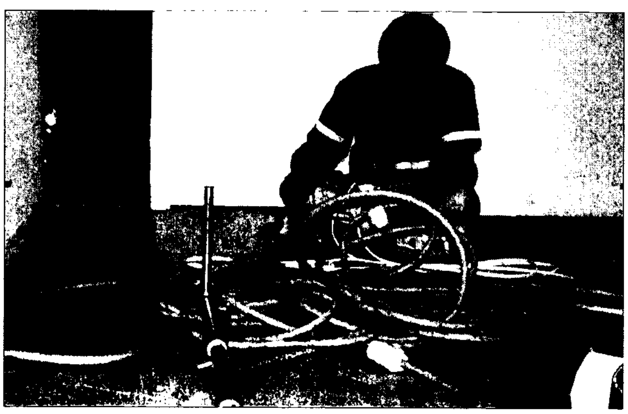

Figure 6: David Henshaw, a friend of the author who just "dropped by" to see the installation of Telematic Connections at the San Francisco Art Institute, who was drafted to splice cables.

Figure 7: "Minds of Concern," 2002, Knowbotic Research. Installation view at the New Museum, New York.

be built today that does not assume that significant network and computing resources will be required in potentially any area of the institution at one time or another.

The Challenge of Legal Bug

A "legal bug" is a concept that the artist group Knowbotic Research coined when their installation for *Open Source Art Hack. Minds of Concern* was shut down not because of infrastructure issues, exactly, but because of a contractual obligation – supposedly – on the part of the institution hosting the show.[4]

Minds of Concern uses port scanning, a technique that is sometimes used by hackers to determine whether there are any weaknesses in a server. Knowbotic's use was simply the scanning, no hacking. After extensive consultation with legal experts around the United States, it was determined that this is essentially like looking in a window or open door from across the street. As long as you don't enter, there is no crime. For Knowbotic's use, they were alerting various non-profits when there was an insecurity in their system. It did not matter that this was a legally protected activity – or at least not illegal activity – the museum's "shrinkwrap" contract with its upstream Internet Service Provider included a blanket clause – apparently standard - that no port scanning was allowed, regardless of intent.

Knowbotic's point was that there is the possibility – and necessity – of a public domain in the digital realm, but that regardless of the public law around the issues, which in itself is problematic, the standard operating practice of shrinkwrap licenses and their equivalents was severely restricting the actual scope of the public domain. Legal bugs, so to speak, are undermining public space in the digital realm.

My point is that institutions, while often overwhelmed by the financial burden of litigation, understand and can protest cogently and strongly an artist's right to fair use, to parody, etc. But in the digital domain, it is often terra incognita, and so much easier to simply say "it's in the contract," and let the lights go dim. As the digital sphere becomes increasingly privatized, interfacing it becomes increasingly compromised.

The Challenge of Presentation

One can list dozens of other challenges to interfacing the digital, but I would like to end with three examples of the presentation of work in physical space. These are not intended, naturally, as universal solutions, but as case studies of attempts to solve particular issues in specific situations.

Let's Entertain and *Art Entertainment Network*

As I said earlier, the interface that we created for the online exhibition *Art Entertainment Network* was designed as a portal; a format "native" to the network. Once we had decided on this exhibition design, we commissioned Antenna Design in New York to create a physical interface which could be used in the galleries as part of the parallel exhibition of visual arts, *Let's Entertain* (http://www.walkerart.org/va/letsentertain/le_content.html).

Antenna designed a freestanding, revolving door, which acted as a kind of portal between the physical space of the exhibition and the virtual space of the online artworks. As you push the door around, it automatically calls up the home page of each project. A touchpad allows you to interact with the work.

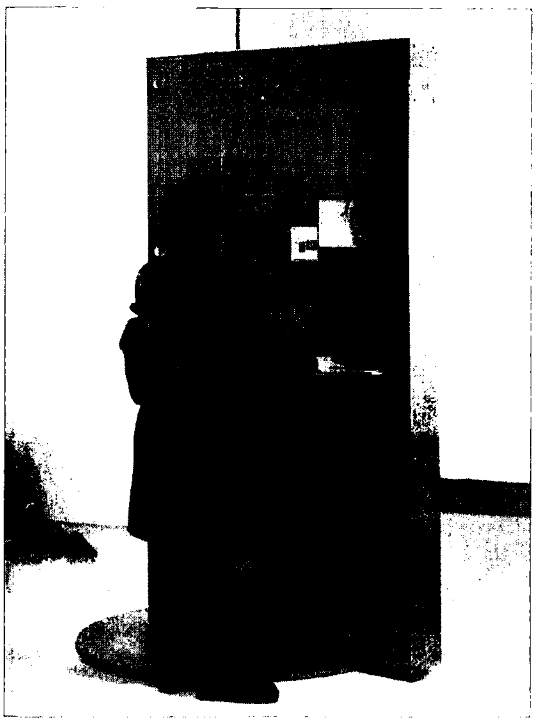

Figure 8: Door-Portal for *Art Entertainment Network* as part of *Let's Entertain*, Walker Art Center

This door could hold its own, so to speak, with the other installations in the exhibition. At the same time, it was appropriate to the concept of the online interface – as a portal. It also didn't assume that the goal of the interface was to create a comfortable browsing situation for hours of enjoyment. Like much gallery behavior it was designed for more casual browsing. A holder next to the didactic label contained printed bookmarks which visitors could take and use to later log on to the site at their convenience and in their favorite viewing position.

Architecture for Temporary Autonomous Sarai

The Walker's most recent commission is a collaboration between Raqs Media Collective from Delhi and Atelier Bow-Wow, an architectural practice in Tokyo, *Architecture for Temporary Autonomous Sarai*, which is part of the *How Latitudes Become Forms* exhibition, currently on view.

From a visit to Raqs Media Collective's Sarai project – first online and then in Delhi – I was very impressed with their genealogy of the idea of the sarai as well as with the energy, levels of interaction, and quality of output. This is how Raqs described Sarai in a conversation with myself, Yukiko Shikata, and Gunalan Nadarajan that undergirded the parallel online exhibition, *Translocations*.

> ". . . for us, the creation of a *sarai* was to create a 'home for nomads' and a resting place for practices of new media nomadism. Traditionally, *sarais* were also nodes in the communications system (horse-mail!) and spaces where theatrical entertainments, music, dervish dancing, and philosophical disputes could all be staged. They were hospitable to a wide variety of journeys—physical, cultural, and intellectual. In medieval Central and South Asia, *sarais* were the typical spaces for a concrete translocality, with their own culture of custodial care, conviviality, and refuge. They also contributed to syncretic languages and ways of being. We

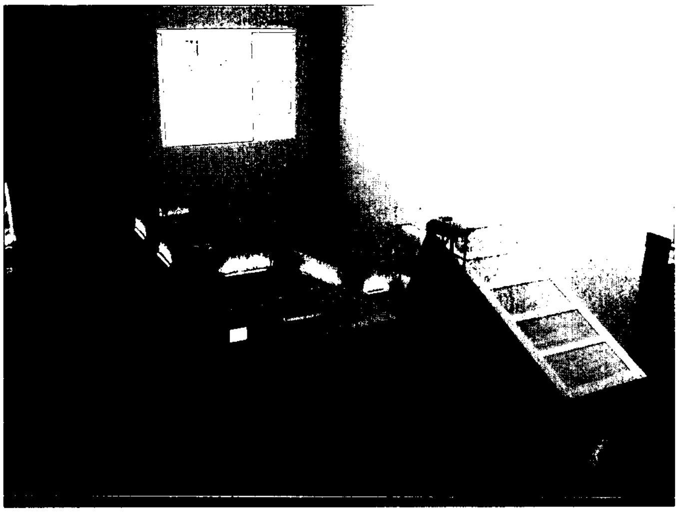

Figure 9: Installation view of Architecture for Temporary Autonomous Sarai by Raqs Media Collective and Atelier Bow-Wow, realized February 2003, Walker Art Center.

would do well to emulate even in part aspects of this tradition in the new media culture of today....This might create oases of locatedness along the global trade routes of new media culture.(Transcript, Translocations, full transcripts at: http://latitudes.walkerart.org/translocations/)

A sarai was exactly what was needed for the *How Latitudes Become Forms* exhibition – a place for social intercourse, both onsite and translocally; a place for the investigation of both artists' work and the exhibition context.

Another artist group in the exhibition was the Tokyo-based architectural practice Atelier Bow-Wow, Yoshiharu Tsukamoto and Momoyo Kaijima, who are proponents of what they have named *da-me* or no-good architecture. Multilayered structures with varied uses (underpass + cinema + bar + barbershop + store, for example), these buildings epitomize, for them, a new creative, adaptive aesthetic.

We decided to ask Raqs and Bow-Wow to collaborate on a "Temporary Autonomous Sarai" – something that was physically modest, intended to be temporary, and programmatically could function as a sarai for the exhibition.

I see the cross-disciplinary collaboration and the physical, social space for presenting net art and its context as an exciting experiment which will inform future practice at the Walker.

Telematic Table

As I said, we believe that artist practice can and should inform institutional practice. One of the projects the Walker has commissioned for the opening of our new spaces is a "telematic table."

The core idea of this interface is to use the format of the table to create a social learning platform that is different than the typical nose-to-screen kiosk experience.

In 2001, after receiving a grant from the NEA, we held an international design competition, inviting over 30 artists, designers, and architects to submit designs for a telematic table with only this for criteria:

Figure 10: Design for Telematic Table by Marek Walczak, Michael McAllister, Jakub Segen, and Peter Kennard; commissioned by the Walker Art Center

> We are envisioning a human-scaled interface that is neither a standard desktop computer nor a public kiosk, but which viscerally engages the user; encourages social interaction among groups of people, can be networked and adapt to a variety of situations and museum spaces. Like an ordinary table, the telematic table is a space of gathering and exchange. It will give its users access to the Walker's multidisciplinary collections and resources, foster curiosity and inquiry into the museum's information assets, and create a setting for social interaction and dialogue among groups of visitors. *(Walker Art Center, 2001)*

Of the responses, we selected 5 for further paper prototyping and from those selected a proposal by a virtual group composed of Marek Walczak, artist/architect; Michael McAllister, furniture designer; Jakub Segen, Bell Labs researcher; Peter Kennard, programmer.

The core idea of the table is very simple. Using gesture recognition software developed by Segen, which allows for a multiple-touch interface, users drag digital assets from a "pond" to personal "puddles." Relationships with related works in other puddles are automatically highlighted, encouraging, hopefully, cross-tabletalk and interaction. In any case, just the shoulder-to-shoulder layout of the table, we expect will lead to at least social situations, if not direct conversations. Naturally, more information is available about each selected object; objects can be "collected" on a postcard printout; and URLs for even deeper investigation from home will be provided.

Conclusion

Even though the Turing Machine – the computer – is defined as a universal black box that can do anything - follow any instruction set — it is important not to confuse and conflate the flexibility of *computation* with the physical interface of computers as they are currently sold by most corporations. Interfacing the digital can and should be as varied as the situations in which it is called for; the audiences being sought; and the artwork being presented. This will take imagination and ingenuity, but it is time for museums to stage their own version of the infamous Apple ad that threw a "monkey wrench" into the Orwellian screen of IBM computing and be as creative interfacing the digital as they are with every other exhibition installation they present.

Notes

1 See http://www.heise.de/tp/english/pop/event_1/4079/1.html for a review of the Documenta X website, which was also "appropriated" by Vuk Cosic and is viewable at http://www.ljudmila.org/~vuk/dx/. The website for the 2000 Whitney Biennial is viewable at http://www.whitney.org/exhibition/2kb/internet.html, and the *net condition* is at http://on1.zkm.de/netCondition.root/netcondition/start/language/default_e.

2 See http://www.nettime.org; http://rhizome.org; http://bbs.thing.net, http://www.jiscmail.ac.uk/lists/NEW-MEDIA-CURATING.html and http://www.archimuse.com/mw98/beyond_interface/.

3 See http://translocations.walkerart.org; specifically http://latitudes.walkerart.org/artists/index.wac?id=80 http://opus.walkerart.org http://latitudes.walkerart.org/artists/index.wac?id=273 http://translationmap.org http://latitudes.walkerart.org/artists/index.wac?id=221 http://latitudes.walkerart.org/translocations/recombo/ and http://latitudes.walkerart.org/artists/index.wac?id=271

4 See http://www.netartcommons.net, http://www.netartcommons.net/article.pl?sid=02/04/26/0311201&mode=thread and http://unitedwehack.ath.cx/barca/barca1/.

References

Walker Art Center, 2001, Proposal for Telematic Table, Submitted to the National Endowment for the Arts (NEA).

Design for Learning

Practicing What We Teach: How Learning Theory Can Guide Development of Online Educational Activities

David T. Schaller and Steven Allison-Bunnell, Educational Web Adventures (Eduweb), USA

Abstract

Since the World Wide Web became in 1994 the first new mass medium since television, online learning design has evolved at Internet speed, taking in less than a decade what it took exhibit design over a century to develop in sophistication. Although virtual exhibits consisting of pictures and text are still common, educational Web designers increasingly employ techniques borrowed from interactive exhibit developers, video game producers, and museum educators to create compelling activities that fully exploit the strengths of the new medium. Constructivist learning theory often informs these new approaches. However, transplanting learning theory from the classroom or museum environment to the Web poses unique challenges. In this paper, we review several theories of learning and explore ways that we have tried to incorporate them into our development and design process for interactive Web sites.

Constructivism underlies much educational practice in museums and is the basis for all of the learning theories we survey in this paper. Each of these, however, clarifies, expands upon, or revises the notion of constructivism in ways that can help Web designers better conceptualize and execute their projects. For example, Kolb's model of learning styles highlights the structure of the learning process. This model offers insight into how to make Web media go beyond the convergent/logical learning that comes easiest to computer-based learning, and to teach divergent, practical, and social learners. Similarly, Gardner's checklist of multiple entry points offers a valuable perspective on diversity in learning, prompting us to look for ways to engage various intelligences in one package. Most dramatically, Egan's notion of developmental "kinds of understanding" frees us from the strict constructivist demand to account for the concrete prior knowledge of our mostly anonymous online audiences. Instead of attempting that impossible feat, or ignoring the issue entirely, we can engage children's and adults' imaginative capacities with stories about profound abstractions, the limits of reality and experience, and our place in the world.

Keywords: Learning Theory, Constructivism, Multimedia Web Development, Online Learning, Evaluation and Research

Introduction

For the purposes of this paper, we wish to offer a fairly demanding definition of "interactivity" as applied to online learning materials. Given the present pervasiveness of random-access devices, from audio CD-players to even the most rudimentary Web sites, merely being able to choose what to see or read is no longer very challenging to conceptualize or produce. Reaching higher, we see the paragon of interactivity involving at least some of the following, non-mutually-exclusive elements:

- Enabling two-way communication between real people, whether in real time or not. Failing that, simulating the richness of two-way person-to-person communication with the computer via narrative devices, visual and interface design, and the underlying information architecture is a common way to automate what is otherwise time and resource-intensive.
- Offering the user "productive decision-making opportunities" (Strohkorb 2002). This means going beyond offering menus of choices to select from, and involves providing an infrastructure for helping the users to construct their own chain of inference and meaning. Giving the users the ability to experience the consequences of their choices is at the heart of state-of-the-art interactivity and game-based learning theory (Prensky 2001, Schank 1992).
- Giving the users something to *do* rather than something to see. Manipulation of objects and ideas that produce a new construct is crucial to a strong form of interactivity (see discussion of "Creative Production" activities below). Merely manipulating an interface widget is not in and of itself interactive in a meaningful way.

Defining what makes a Web site "educational" is just about as slippery as nailing down what counts as interactive. Again, we favor a strong definition that includes these basic characteristics:

- Learning goals or outcomes that can be explicitly articulated by the designers (if only after the fact)

- A clearly focused subject domain

- Scaffolding to help the users develop a skill or increase knowledge under the structured guidance of the program. Thus, a database, by itself, may be a tremendous research tool, but it may not actively help the users develop expertise in using it or understanding the subject domain (Jonassen *et al*, 1999).

The learning theory most commonly known as constructivism holds that learning "is not a simple addition of items into some sort of mental data bank but a transformation of schemas in which the learner plays an active role and which involves making sense out of a range of phenomena" (Hein 1998). Designing learning experiences that facilitate such active meaning making while engaging a generation raised on a steady diet of video games and fast-paced television is the ongoing challenge facing formal and informal educators alike. Computers add to the challenge, since it is all too easy to design computer programs that substitute sophisticated but passive representations for the healthy (and fun) cognitive work of learning. How can we, as practitioners of educational Web design, apply current theories of learning to our development process and final products?

One Size Fits None

While the theoretical and descriptive frameworks of learning theory vary widely, there is consensus that people perceive and process information in many ways. Some of these theories apply general personality theory to learning while others derive from studies of the learning process itself. Each seems coherent enough on its own. But can they all be right? Should we try to accommodate them all in our development philosophy and process? Guild

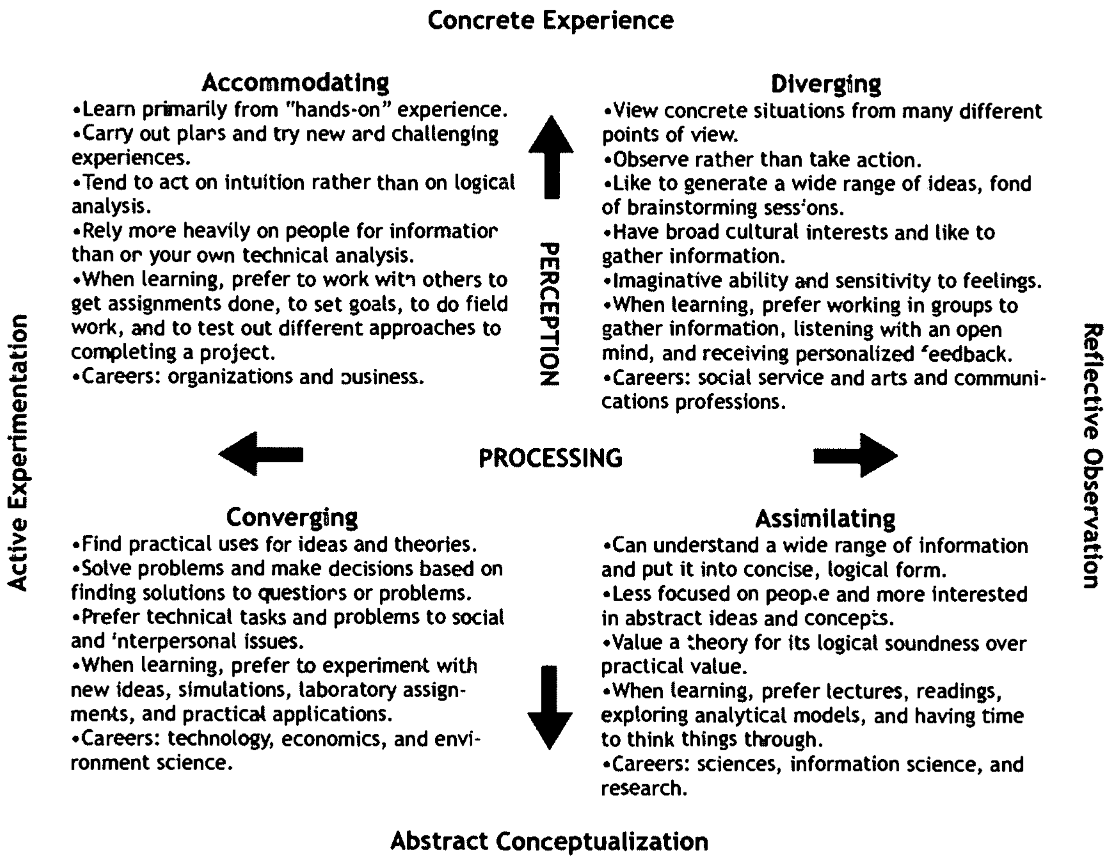

Figure 1. Kolb's model of learning styles. Adapted from Kolb et al. 1999.

and Garger (1998) suggest that we are in a pre-paradigmatic phase of learning style research. The blind researchers have each described a different part of the elephant, but have not yet synthesized their findings into a picture of the whole beast.

Of the many theories about learning styles, we have found David Kolb's experiential learning theory (ELT), first developed in the 1970s, to offer us valuable insights. Kolb emphasizes the importance of experience in the learning process and draws on research by Dewey and Piaget, among others, to identify two major dimensions of learning: perception and processing. Each dimension has two extremes: perception ranges from concrete experience to abstract conceptualization, and processing ranges from reflective observation to active experimentation. These two dimensions form a four-quadrant model (Figure 1).

Kolb's model of learning styles reveals another challenge: how to engage these various ways of learning via a flat computer screen. The Web as an "information superhighway" is ideal for Assimilating Learners who prefer reading and reflective analysis. But what about Accommodating Learners who prefer social environments, or Divergent Learners who need to brainstorm and obtain personal feedback? Can online activities effectively engage them? We are currently planning a research study to investigate these questions further, but we have already made some attempts to put Kolb's theory into practice. *The Artist's Toolkit* (http://www.artsconnected.org/toolkit), developed with The Minneapolis Institute of Art and the Walker Arts Center, offers four ways to explore a series of art elements and principles (Table 1). The first three components were designed for K-5 grade students; the fourth one is a didactic resource for older students and teachers.

Artists' Toolkit Component	Perception-Processing Axis	Learning Style
Watch: Short animations that demonstrate the concept using an artwork from the institutions' collections (Figure 2).	Reflective Observation	**Diverging and Assimilating:** • View concrete situations from many different points of view. • Observe rather than take action. • Less focused on people and more interested in abstract ideas and concepts. • When learning, prefer lectures, readings, exploring analytical models, and having time to think things through.
Find: Examine three artworks and drag labels of each concept over to instances of them in the art (Figure 3).	Active Experimentation	**Accommodating and Converging:** • Learn primarily from "hands-on" experience. • Solve problems and make decisions based on finding solutions to questions or problems. • Prefer technical tasks and problems to social and interpersonal issues. • When learning, prefer to experiment with new ideas, simulations, laboratory assignments, and practical applications.
Create: Apply the concepts in an original composition with an open-ended picture-making tool. Add lines and shapes to a canvas and arrange them to create a picture (Figure 4).	Concrete Experience	**Accommodating and Diverging:** • Learn primarily from "hands-on" experience. • Tend to act on intuition rather than on logical analysis • Like to generate a wide range of ideas, fond of brainstorming sessions. • Imaginative ability and sensitivity to feelings.
Encyclopedia: Study the concepts in more depth through expository text and example artworks (Figure 5).	Abstract Conceptualization	**Assimilating:** • Can understand a wide range of information and put it into concise, logical form. • Less focused on people and more interested in abstract ideas and concepts. • Value a theory for its logical soundness over practical value. • When learning, prefer lectures, readings, exploring analytical models, and having time to think things through.

Table 1. The Artists' Toolkit components and targeted learning styles.

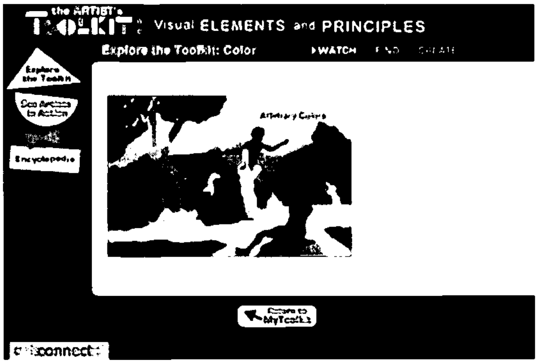

Figure 2. Watch module. www.artsconnected.org/toolkit/watch_color_arbitrary.cfm

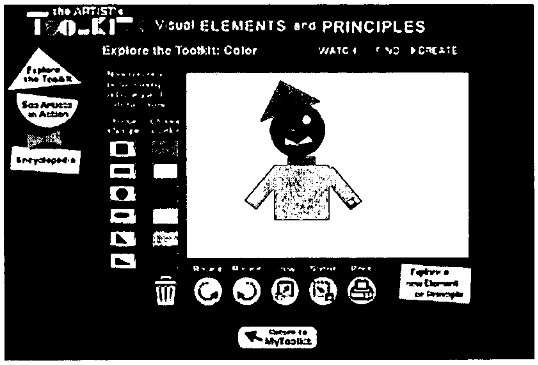

Figure 3. Find module. www.artsconnected.org/toolkit/find_color_arbitrary.cfm

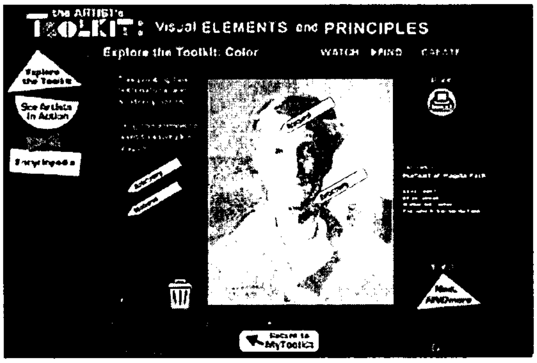

Figure 4. Create module. www.artsconnected.org/toolkit/create_color_arbitrary.cfm

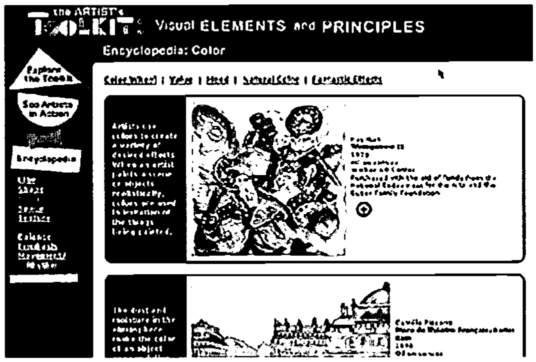

Figure 5. Encyclopedia module www.artsconnected.org/toolkit/encyc_colornatural.html

Connecting Learning Styles and Activity Types

By examining differences in mental processing preferences, Kolb's model also suggests ways that online activities can engage different types of learners. We have hypothesized plausible links between Kolb's styles and our own typology of online activities (Schaller *et al.*, 2002) as follows:

Role-Play activities allow users to adopt a persona different from their own, giving them the ability to do things they cannot ordinarily do (e.g. break natural or societal laws, experience people and places normally out of reach). They can also interact with other characters, whose behavior either may be scripted or controlled by other players. (True social interaction on the Web, i.e. direct communication with other people, is one of its key strengths, but it also poses many challenges with a school-age audience, from privacy concerns to risks of inappropriate communication. This means that many activities that could in principle take advantage of the technology to connect real people with one another must substitute interaction with the program.) We hypothesized that the Role-Play activity type may appeal to those with a strong Accommodating learning style.

Simulation employs a model of the real world that users can manipulate to explore a system. It involves direct engagement with representations of data with some degree of generalization or abstraction. While many activities may have an underlying simulation or model that generates the activity decision tree, "simulation" as used here is explicitly framed as such in its presentation to the user, and is used to develop conceptual understanding of a complex system. We hypothesized that this activity type may appeal more strongly to people with either an Assimilating or Converging learning style.

Figure 6. Scene from Squish the Fish. www.sheddaquarium.org/SEA/ interactive_module.cfm?id=8

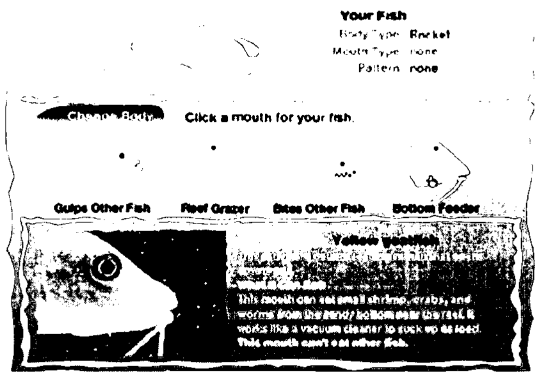

Figure 7. Mouth-types choice screen from Build-a-Fish www.sheddaquarium.org/SEA/ interactive_module.cfm?id=7

Puzzle/Mystery involves analysis and deductive or inductive reasoning to reach a logical solution. The user relies on evidence from people, nature, or reference material provided by the activity to solve the problem. We hypothesized that this activity type may be more appealing to people with either an Assimilating or Converging learning style.

Creative Production emphasizes both open-ended self-expression and the application of subject knowledge and concepts to building some kind of product such as a story, a picture, a movie, etc. We hypothesized that this activity type may appeal more often than the other types to the Divergent learning style.

With these associations in mind, we sought to offer a variety of learning experiences in a suite of educational Web activities, *Shedd Educational Adventures (SEA)* (www.sheddaquarium.org/sea), developed in partnership with the John. G. Shedd Aquarium. We should emphasize that we do not see each activity type appealing solely to a single learning style. Rather, different learners will utilize each type of activity differently. As developers, recognizing the strengths of each activity type and the learners it appeals to most strongly helps us avoid building activities that appeal merely to our *own* learning styles.

Squish the Fish (Figure 6) is an interactive role-play story for primary-grade children (grades K-2). Although users do not explicitly assume the role of Squish, they guide him through a dangerous journey across a coral reef. Squish has the ability to shape-shift and color-shift to avoid predators, and users can vicariously go along for the ride. While experiencing no social interaction with other people in real time, users do meet a variety of talking fish and other undersea creatures that personify the various adaptations that Squish uses to stay alive, thus attempting to engage Accommodating learners.

Build-a-Fish (Figure 7) is a simulation with a touch of creative production, intended for upper primary students (grades 3-5). Players design a fish and then navigate it around a detailed coral reef, trying to eat other fish and avoid being eaten in turn. This pragmatic challenge engages Convergent Learners, while Assimilating Learners can review didactic content to inform their hypotheses of which fish designs will be most successful.

Mysteries of Apo Island (Figure 8) is designed for middle school students (grades 6-8), and offers clues that the user must analyze and synthesize to solve a series of mysteries. While exploring a map of Apo Island in the Philippines, users encounter stories and personal observations of strange animal behavior. Each of these can be collected in the users' notebooks, where they can then compare the clues and sort them into groups by species. In this way the activity is an "open-ended mystery." While the clues do match up to correct solutions, users are not funneled toward solutions, as they would be in a branching story structure. Instead, users must use logical deduction to hypothesize solutions and then ask a local biologist to check their conclusions.

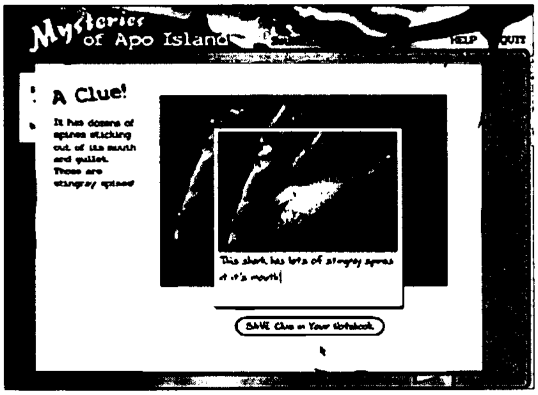

Figure 8. Clue screen from Mysteries of Apo Island. www.sheddaquarium.org/SEA/ interactive_module.cfm?id=9

Figure 9. Seahorse video clip from Conservation Investigation: Seahorses. www.sheddaquarium.org/SEA/ interactive_module.cfm?id=11

We initially planned to develop a simulation for the high school audience (grades 9-10). But due to conceptual issues (detailed below) and budgetary constraints, we turned instead to an interactive reference module. In *Conservation Investigation: Seahorses* (Figure 9), students can explore a conservation issue through primary and secondary sources and write a newspaper editorial based on their research. The students assume the role of investigative reporters and are offered a range of materials about seahorse conservation grouped according to how journalists might research the story: field observations, interviews, and library materials. In writing the final product, the students are guided by a template that encourages logical organization and solid evidence to supporttheir opinion. This combines role-play and creative production elements, but in a less fanciful, more serious fashion for this more mature age group.

Integrating Learning Styles into One Activity

One of the challenges of using Kolb's framework is that while it is clear how to match a particular learning style with a particular activity type, Kolb offers less guidance in creating a single activity that can meet the needs of learners with multiple learning styles. Howard Gardner's theory of multiple intelligences has earned wide acclaim and regard from both formal and informal educators. By examining many cognitive modalities, his theory inspires us to stretch the Web's capabilities to engage many different kinds of learners. Table 2 shows how Gardner's Entry Point Approach outlines ways to provide multiple entry points to a single learning experience "to accommodate the different lenses through which learners see" (Davis and Gardner, 1993).

Entry Point	Learners respond to
Narrative	The story.
Numerical	Numbers, statistics, and mathematical models, as well as musical rhythm.
Logical	Logical propositions, syllogisms, narrative plot structure and cause and effect relationships.
Existential/ Foundational	Big questions about life, death and our place in the world.
Aesthetic	Surface qualities of art and music, as well as other subjects, i.e. balance in an ecosystem, harmony in the built environment, etc.
Hands-on or Experiential	Direct exploration, both physical and virtual, through experimentation, creativity, and immersive experiences.
Interpersonal	Social interaction through cooperation, debate, and role-play.

Table 2. Howard Gardner's Entry Points for multiple intelligences (Gardner 1999).

Several researchers have explored these entry points in the museum context. Project Explore examined which entry points children and adults utilized at several children's museums, and found that Experiential was by far the most common, with Narrative a distant second and the rest even less frequently evoked (Please Touch Museum and

Activity	Entry Points	How users might engage entry points
Squish the Fish	Narrative	Responds to Squish's quest to travel across the reef to eat and to return home again.
	Aesthetic	Comparing examples of shape, color, and pattern adaptations. Reading the rhyming verse.
	Hands-on/ Experiential	Trying an adaptation and seeing whether it protects Squish from Big Tooth Blob.
	Logical	Systematically testing each adaptation on various places on the reef.
Build-a-Fish	Hands-on/ Experiential	Designing a fish and seeing whether it can eat and avoid being eaten on the reef.
	Logical	Systematically testing fish designs to see which does best in various places on} the reef.
	Aesthetic	Considering shape and color differences in design options while designing fish.
Mysteries of Apo Island	Narrative	Responds to the story setup.
	Logical	Examining and comparing clues, identifying shared characteristics of clues. Comparing solutions to biologist's fact cards.
	Hands-on/ Experiential	Exploring the island map, sorting clues into groups.
Conservation Investigation: Seahorses	Narrative	Responding to journalist story framework.
	Logical	Recording, reviewing and analyzing research information and synthesizing into article.
	Interpersonal	Observing video clips of seahorses and interviews with biologists.

Table 3. Connections between Entry Points and Shedd SEA activities

Project Zero, 1998). Experiential is, of course, the easiest entry point to observe in a museum setting. But even in an exhibit with musical instruments, the Aesthetic entry point was rarely observed. The researchers concluded that this disparity may stem from either methodological issues or deeper questions of exhibit design or even young children's cognition.

For Shedd's SEA package, we considered how to design each online activity to contain multiple entry points (Table 3).

During pilot testing, naturalistically observed student behavior suggested that the entry points we intended for the activities were in fact used.

Squish. The target audience responded enthusiastically to the Narrative framework of Squish.

Build-a-Fish. Younger students (3rd grade) typically used the Hands-on entry point, experimenting with fish designs without logically considering the causes and effects of each choice. Middle school students, however, usually did prefer the Logical entry point and were often quite methodical in testing each combination of fish design options. Aesthetic engagement was seen in requests for a larger variety of fish body parts and colors to choose from. Further evaluation is required to determine whether the game was simply too challenging for third graders to use more logical approaches without teacher guidance or whether children favor different learning styles at different developmental stages.

Mysteries of Apo Island. Middle school students responded well to the Narrative entry point, sometimes making up their own scenarios about why a shark might behave in a way consistent with the clues. As is discussed below, younger middle school students had trouble with the Logical aspects of the activity, and frequently were unable to develop the chain of inference that would let them solve the mystery.

Conservation Investigation: Seahorses. Teachers evaluating this activity felt that the Logical guidance offered in structuring the research and writing exercise would be helpful and appropriate for this age level.

Kinds of Understanding

As noted above, current learning theory holds that learning involves a transformation of mental schemas, from flawed and incomplete to truer and more complete (Bransford *et al.* 2000). To facilitate such transformations most effectively, educators must understand their students' prior knowledge,

Kinds of Understanding	Age	Concerns	Examples
Somatic	Birth to three	Body abilities	Walking
Mythic	Three to eight	Binary opposites	*Star Wars* movies
Romantic	Eight to fifteen	Limits of reality and experience; heroes, idealism	*Guinness Book of World Records* Michael Jordan "Save the Earth"
Philosophic*	Fifteen to twenty	Systems and schema	Marxism
Ironic*	Twenty and up	Self-reflective	Post-modernism

Table 4. Egan's theory of developmental phases in learning style (Egan 1998).
*While the first three phases occur fairly naturally in human development, the latter two require substantial guidance and support for an individual to achieve.

experiences, and most critically, their misconceptions about the subject. Only by tackling such misconceptions directly can educators effectively help learners to overcome them and achieve a truer (more accurate) understanding of the world (Borun 1990, Hein 1998). Much of our knowledge of such misconceptions derives from Piaget's research. Indeed, much of contemporary learning theory is based on his work, which mainly investigated children's learning about the physical world and mathematics (Egan 1998). In such domains, hands-on experimentation is invaluable as a way to bring "students' earlier models or misconceptions into sharp focus [through] an experience that directly challenges the viability of the model they have been favoring" (Gardner 1991). Exploring a parallel phenomenon in the humanities, Gardner discusses the persistence of stereotypes, scripts and simplifications such as "the 'bad man' theory" of history, even among well-educated college students (*Ibid.*). To transform such misconceptions, a strict constructivist approach would start with familiar and concrete objects and places before moving to the distant and abstract.

A teacher can discover her students' misconceptions and stereotypes from a class discussion at the start of a lesson. Similarly, museums conduct front-end research studies to ascertain their audiences' pre-existing knowledge and misconceptions and stereotypes (also called preconceptions and naïve notions) about a program or exhibit topic. This task is challenging enough when exhibits attract visitors of many ages and experiences. On the Web, which attracts learners of all ages in varying settings from around the globe, trying to determine pre-existing knowledge and misconceptions is even more daunting. When learning theory calls for an emphasis on what is familiar to learners, the challenge becomes difficult indeed.

Educational theorist Kieran Egan offers a way out of this dilemma by re-imagining what *kinds* of knowledge come naturally to learners at each stage of development (Egan 1988, 1992, 1998). He argues that our powerful imaginative capacities are much more effective (and *affective*) avenues to learning than the prosaic approaches advocated by Piaget's strictest adherents. Even young children have a deep (if difficult to articulate) understanding of certain profound abstract concepts such as good/bad, beauty/ugliness and survival/destruction. Medieval fairy tales and the *Star Wars* movies would have little hold on children if this were not the case. These abstractions have deep affective meaning for children (and adults) and thus are very powerful ways to present new material.

Each new kind of understanding builds on previous stages. As young children, we begin to organize the world by dividing everything into binary opposites. Only after we have created basic categories can we start to refine our mental organization of the world, typically by exploring the extremes of reality (from the tallest human to a hero who transcends limitations and adversity) to better understand our own place within this order. It is after gaining a sense of the range and variety of experiences that we start thinking about how they all relate to one another. The philosophic level propels the search for causal connections that have birthed theories from Marxism to the second gunman on the grassy knoll. Acquiring the ability to construct such general schema can trigger a period as a "true believer" with tunnel vision, but with some outside support, we can ultimately achieve Ironic understanding, or a deep appreciation that every theory contains some truth but none can claim to be The Truth. This doesn't mean that adults who have reached this phase can't enjoy an occasional *Star Wars* movie, for the "simplistic" theme of good vs. evil still resonates deep within us, but that we can appreciate it for what it is without then interpreting real international affairs through the same lens.

What Egan's theory offers us as designers of online learning experiences is valuable guidance about the kinds of abstractions people will find innately relevant and meaningful. We were inspired by these ideas during the development of the Shedd Educational Adventures online activities with the Shedd Aquarium. Since this project entailed producing online learning modules for five grade groups: K-2, 3-5, 6-8, 9-10, and 11-12, we kept Egan's kinds of understanding in mind during our development process, letting them inform the themes we developed and the stories we told.

Mythic Understanding and Squish the Fish

The learning goals for elementary grade activities centered on fish adaptations and easily lent themselves to stories of danger and survival. Squish the Fish and Big Tooth Blob, cartoon characters from Shedd's existing education materials, embody the struggle for survival which children so easily grasp. Squish must travel across the reef, looking for friends who have successful shape, color and behavioral adaptations to protect themselves from Big Tooth Blob. Children instinctively understand the binary opposites of safety and danger and this, rather than any reference to the concrete world of our child-audience, is the theme that carries the message.

Amplifying the safety-danger theme meant making Big Tooth Blob a pure predator with no sympathetic traits. However, since his name and appearance are more comical than threatening, he is portrayed as simply hungry, rather than as menacing or malevolent. Thus no one clamored to give Big Tooth's side of the story. With a subject such as history, resistance to binary stereotyping from either the target audience or teachers and subject experts would pose more problems.

The real question is whether the story is mythic *enough*. Using the term in a similar way, Rounds (2002) describes the tension between mythic and scientific thinking, arguing that museum visitors often interpret science content in mythic terms. The factual content is not what interests them. It is the meaning they can make from the stories that is important, and they judge those stories by their affective authenticity, not by their factual accuracy. The mythical aspect of the safety-danger theme in Squish gives the story some degree of affective authenticity, but our real goal centered on communicating the factual content about various organisms and their adaptations. Squish succeeds because it offers a scientifically credible mythic framework for the more abstract and emotionally neutral concept of adaptation as an evolutionary response to survival and predation. We want to emphasize that deploying a mythic theme that does not naturally stem from the content in an activity such as this will only undermine its educational value for the learner and its credibility with the teacher.

Romantic Understanding and The Mysteries of Apo Island

The Shedd's learning goal for this activity was shark biodiversity—showcasing "the many kinds of sharks and their fascinating adaptations and behavior." Exploring this theme through the romantic fascination with the exotic and bizarre ("the limits of reality and extremes of experience," in Egan's terms) was straightforward enough, since there are plenty of odd shark adaptations and behaviors from which to choose. The challenge lay in taking it beyond a "trivia quiz" approach that might prove intriguing but would still consist of informational dead-ends. Our goal was to design an activity that leveraged the romantic fascination with exotica into a larger learning experience that developed synthetic reasoning skills. For that we turned to a mystery format. The trivia were made into clues that, if examined and integrated, provided the solutions to the mystery. Thus we hoped to foster a romantic understanding of shark biodiversity within a deductive reasoning process, thereby suggesting the path ahead to philosophic understanding.

Pilot testing indicated that shark exotica were indeed successful in engaging the students' attention. Some also responded quite enthusiastically to the deductive structure of the activity, while others (especially students at the lower end of the grade 6-8 target range) enjoyed the clues but found the mystery befuddling. That kind of synthesis is more characteristic of philosophic understanding, and so it shouldn't be surprising that some 12-year-olds found it too challenging.

In another example of Rounds' mythic-scientific distinction, we also found that students sometimes assumed the solution would focus on an *individual* shark rather than a shark species as a *group*. Students would spin their own stories about "a bad shark" that attacked the boat and other sharks, fantasizing many details that were quite out of character for the mystery itself. Of course, the mystery format (as well as most appearances of sharks in popular culture) encourages this interpretation. Our own scientific-mindedness had blinded us to the likelihood of this result. We have since emphasized the idea of species as groups in the activity text, but we await summative evaluation to see whether that makes any difference.

Philosophic and Ironic Understanding and an Ill-fated Simulation

Inspired by Sherry Turkle's call for "simulations that teach about the nature of simulation itself, that teach enough about how to build [your] own simulation that [you become] a literate 'reader' of the new medium," we decided to tackle just such an effort (Turkle 1997). Building on a coral reef management simulation already produced for Shedd, we imagined letting students "under the hood" of the simulation so they could tinker with its assumptions and parameters. We were motivated by our own awareness that every model has both descriptive power and predictive blind spots. We wanted students to be able explore how the hidden assumptions and underlying logic of a model can dramatically affect the simulation's portrayal of the real world.

However, presenting this concept to high school teachers in focus groups led us to abandon it. The teachers felt that orienting students in this age group to a model of ecosystem management, however simplified, was challenging enough. Layering additional complexity with the reflexive analysis of the model itself was more than students could handle, at least with such an unfamiliar subject. Consequently, we switched to a simpler and more traditional interactive reference format (*Conservation Investigation: Seahorses*), hoping to return to Turkle's challenge in the future. In retrospect, we should have heeded Egan earlier, since his notions of philosophic and ironic understanding anticipate this quandary. Philosophic understanding, as in looking for systems, schema, and underlying causes, is well suited to explicitly exploring models and simulations as such (rather than concealing the model in a mythic or romantic wrapper). But because philosophic understanding so enthusiastically embraces the search for cause and effect, there is a strong tendency at this stage to take the model at face value. The ability to step back and question the construction of the model itself truly arrives only with Ironic understanding.

Since questioning models is ultimately such an important skill, we still hope in the future to develop a simulation that facilitates this sort of meta-examination. However, it is clear that the simulation should focus on a familiar domain and be targeted at more advanced (perhaps even college-level) students. Our ecosystem management simulation tried to tackle too much and over-ambitiously tried to work on two very different levels: the issues surrounding coral reef management efforts, and the art of modeling a real world system. That is simply too much of an analytical load for high school students to handle.

Offering a simulation as an appropriate means of examining a system's function might also leverage Mythic understanding, capturing the users' interest with a dramatic binary opposition between stereotyped forces of good and evil (perhaps polluters and conservationists), then suggesting that this model does not show what's *really* going on and inviting them to dig under the surface of the simulation in search of the truth. In this way, users would find some initial comfort in mythic understanding while engaging in their natural philosophic quest for truth and causality. Similar approaches could easily be applied to history and other subjects where mythic stereotypes are common, but where Ironic understanding might lend a valuable level of sophistication.

Where Do We Go From Here?

The above examples indicate the potential for using learning theory to guide thoughtful and innovative educational design on the Web. When the Web was in its infancy as a medium, trial and error development was appropriate. No one knew what might work, and rapid changes in the technology created a quickly moving target. As the Web has matured, the need for more explicit best practices and metrics for evaluating the quality and success of our projects has become apparent. A face-to-face learning environment allows the effective teacher to respond to subtle cues about student knowledge, interest, and ability. We do not have that luxury online, and instead we must attempt to formalize within the programming of the activity much of the tacit feedback the teacher can draw upon. As we have suggested above, a development process informed by learning theory helps embed the teacher's expertise into an online learning experience that has no two-way communication between real people.

Even as these case studies demonstrate the benefits of explicitly applying learning theory to online activities, they also raise questions. Can rigorous research and evaluation of online learning activities help us refine learning theory as a whole? Can we empirically demonstrate whether these ideas really do improve users' experience, or even their learning (whatever that ultimately is)? Given the present politics of public education in the United States, which emphasizes high stakes testing and demonstrating student mastery of specific standards, we do not serve our clients or their funders well if we continue to offer only impressionistic accounts of the impact of our projects. Such a promise of more objective evaluation might seem hubristic, given how difficult it is to measure the users' experience in the offline world. However, we anticipate that the same technical limitations of online learning that require us to formalize the content and learning methods will also help us to capture and analyze information that will shed light on the user experience.

Furthermore, can we continue to ignore the setting in which online activities are used, or do we have to pay closer attention to where the users are and why they are using the activity? For example, do we need more support for teachers to accompany each activity intended for classroom use? In the case of SEA, Shedd saw teacher support materials as crucial from the outset. We hope to see with further study how support materials affect the effectiveness of the SEA activities. It may be that teachers view electronic activities as by nature self-contained. Given that the Web allows access to all comers, the challenge remains to offer novice users the right entry point without assuming that they can or will use external supporting materials or guidance.

As we have surveyed learning theory, from Kolb to Gardner to Egan, and tried to make sense of our real-world products in their light, it is clear that no one theory can be our single guiding light. It seems more pragmatic to treat the range of learning theory frameworks as a developer's toolbox from which we can pick and choose according to the needs of a particular project. Given the diversity of learners and learning environments, we favor the instrumental view that we should not attempt to impose a single framework on all subject content, learning goals, use context, and institutional missions and cultures. What works for the context and content of one project may be completely unsuitable for another.

At the end of the day, our thumbnail sketch of learning theory does expose limitations in applying general theories to complex, practical problems. As we discussed above, a strong social constructivist approach to pedagogy that puts a high premium on social interaction and assessing and adjusting to student knowledge may find many current online activities lacking. At the same time, we remain optimistic that some types of online learning experiences open doors to understanding that cannot be offered in any classroom (Prensky 2001). This is not a dismissal of the classroom experience, but a recognition of the ways that computer-aided research tools have dramatically transformed the practice of many research fields, from visualization in the sciences to databases in the humanities. Much of our effort is motivated by the desire to bring these tools within reach of non-experts.

Finally, and most ambitiously, can studying online learning actually advance our understanding of learning, or merely confirm or fail to support existing principles developed for other learning arenas? As the methods and instruments for evaluating online learning mature, we expect online learning to offer a lively test bed for the continuing quest for knowing how we know.

Acknowledgements

The authors wish to thank Minda Borun, Theresa Esterlund, Eric Gyllenhaal, Nancy Ross, and Kris Wetterlund for their insightful comments on earlier drafts of this paper.

References

Bransford, J.D., A. L. Brown, and R.R. Cocking, eds. (2000). *How People Learn: Brain, mind, experience, and school.* Washington, D.C.: National Academy Press.

Davis, J., and H. Gardner. (1993). Open windows, open doors. *Museum News*. January/February 1993.

Egan, K. (1998). *The Educated Mind: How cognitive tools shape our understanding*. Chicago: University of Chicago Press.

Egan, K. (1992). *Imagination in Teaching and Learning:The middle school years*. Chicago: University of Chicago Press.

Egan, K. (1986). *Teaching as Storytelling:An alternative approach to teaching and curriculum in the elementary school.* Chicago: University of Chicago Press.

Gardner, H., (1991). *The Unschooled Mind.* New York: Basic Books.

Gardner, H., (1999). *The Disciplined Mind.* New York: PenguinBooks.

Gyllenhaal, E. D., L. Beaumont, & A. Tyree (2003). Formative evaluation of the On-line Teacher Resources project: Final report. Unpublished manuscript, John G. Shedd Aquarium, Chicago, IL. Available at http://www.sheddaquarium.org/SEA/research/.

Guild, P.B., and S. Garger. (1999). *Marching to Different Drummers,* 2nd edition. Alexandria: Association for Supervision and Curriculum Development.

Hein, G. E. (1998). *Learning in the Museum.* London: Routledge.

Henke, H. (1997). Applying learning theory to computer based training and web-based instruction, last updated April 16, 1997, consulted January 14, 2003, http://scis.nova.edu/~henkeh/story2.htm.

Jonassen, D.H., K.L. Peck, and B.G. Wilson. (1999). *Learning with Technology:A constructivist perspective*. New Jersey: Merrill.

Kolb, D.A. (1984). *Experiential Learning: Experience as the source of learning and development.* New Jersey: Prentice-Hall.

Kolb, D.A., R. E. Boyatzis, and C. Mainemelis. (1999). Experiential learning theory: Previous research and new directions. In R. J. Sternberg and L. F. Zhang (Eds.), *Perspectives on Cognitive, Learning, and Thinking styles.* NJ: Lawrence Erlbaum, 2000.

Prensky, M. (2001). *Digital Game-Based Learning* New York: McGraw-Hill.

Project Zero. (2000). MUSE QUESTs (Questions for Understanding, Exploring, Seeing and Thinking). Cambridge: Harvard University.

Rounds, J. (2002). Storytelling in science exhibits. *The Exhibitionist,* Vol. 21, No. 2 (Fall 2002).

Schaller, D.T., S. Allison-Bunnell, M. Borun, M. Chambers. (2002). How do you like to learn? Comparing User Preferences and Visit Length of Educational Web Sites. In D. Bearman and J. Trant (Eds.), *Museums & the Web 2002: Selected Papers from an International Conference.* Pittsburgh: Archives and Museum Informatics.

Schank, R. (1992). Goal-Based Scenarios. Chicago: Northwestern University Institute for the Learning Sciences. Technical Report #36. Available at: http://cogprints.soton.ac.uk/documents/disk0/00/00/06/24/cog00000624-00/V11ANSEK.html

Strohkorb, J. (2002). Personal communication.

Turkle, S. (1997). Seeing Through Computers: Education in a Culture of Simulation. American Prospect. Vol. 8, Issue 31. March 1 - April 1 1997. Consulted January 14, 2003. http://www.prospect.org/print/V8/31/turkle-s.html.

Veenema, S. and H. Gardner. (1996). Multimedia and multiple intelligences. *American Prospect,* Vol. 7, Issue 29. November 1 - December 1 1996. Available at: http://www.prospect.org/print/V7/29/veenema-s.html.

Evaluating the Authenticity of Egyptian Cartonnage Fragments: Educational Outreach in Search of the Truth

Paul Marty, School of Information Studies, Florida State University; Kim Sheahan, and Ann Lacy, Spurlock Museum, University of Illinois at Urbana-Champaign, USA

Abstract

This paper presents preliminary results from a five year longitudinal study on the efficacy of integrating museum resources into elementary and middle school curricula through educational outreach activities, over the Internet and in the classroom. From 1997 to 2003, museum educators at the University of Illinois' Spurlock Museum have explored how students unable to travel to the museum could make the most effective use of the museum's online resources. Using problem-based learning methodologies, museum educators worked with several local school teachers to develop a series of projects that were closely integrated with the students' classroom activities. This paper evaluates the results of the program and presents the process of developing one of these projects: an analysis of the authenticity of ancient Egyptian cartonnage fragments.

Keywords: evaluation, museum education, middle school, authenticity, forgery

Introduction

This paper discusses the lessons learned from a museum/school educational outreach activity conducted over the past five years at the Spurlock Museum at the University of Illinois. From December 1997 to January 2003, over 300 students in ten different classrooms at seven different schools in Central Illinois have been engaged in an exploration of ancient Egyptian history enhanced by the addition of virtual and physical museum resources. Presented with alleged fragments of Egyptian cartonnage (waste papyrus or linen soaked in plaster, often used to create mummy casings), students – working in teams – were asked to assume the roles of museum curators and research the culture, beliefs, and practices of the ancient Egyptians to determine the authenticity of their artifacts. To facilitate the research process, students had access to online resources specifically developed by museum staff members for this project. The availability of these virtual resources, combined with the presence of the physical artifacts, successfully created an enhanced learning experience for the students.

This paper presents the preliminary findings of this study and documents the methods developed by the museum staff to work with the students. It discusses such issues as the relative benefits of sing physical museum artifacts, the importance of visiting the students' classrooms in person, the difficulties of communicating with students through email and online discussion boards, and the significance of developing online databases and electronic resource guides specially geared for the project. It is our hope that the results of this study will shed new light on how museum educators can work with elementary and middle school teachers to integrate museum resources into their curricula in the classroom and over the Internet.

Background and Significance of Project

In November of 1997, a sixth-grade teacher from a local middle school approached staff members at the Spurlock Museum to see if they would be willing to collaborate on designing a problem-based learning experience (Boud & Feletti, 1991) for her students, who were going to be studying ancient Egypt in December. She wanted her students to assume the roles of museum curators, and then to deal with a real-life problem: the authenticity of one of their Egyptian artifacts is being challenged. Their tasks would be to research the artifact, investigate its authenticity, and determine its provenance, while paying particular attention to repatriation issues (such as the UNESCO 1971 treaty). She wondered if the museum could help her out.

In 1997, the Spurlock Museum had just begun a massive project to inventory, pack, and move its collections of 45,000 artifacts across campus to a new facility (Marty, 2000). As part of this move, the museum was to be closed to the public for over two years, and it was important to the museum staff that they find alternate methods of keeping the museum an active part of the local educational community. A project such as this one seemed like the perfect opportunity to explore the museum's potential role as something more than a field trip destination for area school children. In December 1997, therefore, museum staff worked with this sixth-grade teacher to design a prototype version of what would eventually become known as "Museum Problems in Today's World: Egyptian Mummification." The project was a success, and museum staff members soon worked to revise the program, making it more consistent and less focused on problem-based learning, so that it could be opened to more teachers and more widely distributed among area schools. From December 1997 to January 2003, this program has been run ten times at seven different schools and has reached over 300 elementary and middle school students in central Illinois. When the program began, it was decided that the most appropriate artifact for the exercise would be a piece of ancient Egyptian cartonnage: waste papyrus or linen, soaked in plaster (similar to papier-mâché). Although the museum had several such pieces in its collection, it would have been inappropriate for the museum to send these artifacts to the schools. Each time the "cartonnage project" is run, therefore, reproductions of ancient Egyptian cartonnage fragments—manufactured by museum staff members—are delivered to the schools (see Figure 1).

In addition to these physical artifacts, students are provided with access to the museum's online database systems so that they can access database records for similar pieces, a resource guide to ancient Egyptian mummification specifically created by museum staff members for this project (http://www.spurlock.uiuc.edu/education/resources/mummification), and a series of research documents written by the museum's educators designed to provide guidance for the students as they research the cartonnage fragment. The students also have access to materials in their own school and local libraries and on the open Internet. The students generally work in teams of four or five, usually dividing up tasks so that each student could research a different aspect of the artifacts. Depending on the teachers' schedule, the students have (on average) two to three weeks to use these resources to prepare their findings on the artifacts. At the end of this time, the students make a presentation to a panel of museum staff members, where they announce their findings and make their recommendations.

Figure 1: Sample Cartonnage Fragments

The cartonnage project has been a great success for museum staff, teachers, and students alike. Museum staff members find it an effective way of keeping the Spurlock Museum an active part of the educational community (even when the museum was closed to the public). It provides new opportunities for museum-school interactions, both in the classroom and online. Teachers use the cartonnage project to introduce their units on Ancient Egypt. One of the more popular times to run the program is between Thanksgiving and Winter break, as the highly interactive nature of the project keeps students involved at a time when it is often difficult to keep students interested in school. The students find it an exciting and involving way of exploring ancient Egypt. They feel that they have a personal stake in the project, and some students – months after the completion of the project – have even contacted the museum's educators to ask them whether they have implemented the students' recommendations.

That this project should have been successful is perhaps not surprising. Educational outreach projects connecting museums and schools have received a great deal of attention over the past few years, much of it at the Museums and the Web conference. Faculty from Museum Studies departments as well as museum curators have discussed the pedagogical impact of virtual museums from both educational and museological perspectives (Teather & Wilhelm, 1999; Sumption, 2001). Researchers at the University of Michigan have explored the challenges of coordinating teachers, museum professionals, and information specialists to create online exhibitions (Frost, 1999; Frost, 2001). Researchers at the University of Illinois have explored methods of integrating digital primary source materials from cultural heritage museums into classrooms (Bennet & Trofanenko, 2002; Bennet & Sandore, 2001). Museum educators at the Seattle Art Museum have even used the Internet to help sixth-grade students in Seattle better understand the curatorial process (Adams, et al., 2001). What is there that makes the project at the Spurlock Museum unique?

We feel that the cartonnage project is important for two reasons. First, the project has evolved (and is still evolving) over a period of five years. There is a tremendous amount of longitudinal data available that document what we learned during this time about what worked and why. We can use these data to examine how this project changed over time in response to the reactions, interactions, and levels of involvement of students, teachers, and museum personnel. We can discuss how changes to this project over time affected the students' ability to learn, what they learned, and how they worked within the program. This analysis can help us determine what drove the historical evolution of the project and how changes in the relationships between students, teachers, and museum professionals affect the development of educational outreach projects.

Second, preliminary data analysis of this project revealed an extremely interesting fact: of the 300 students involved in this project, approximately 75% reached the conclusion that the artifacts were authentic. Of the other 25%, the majority were undecided; of those students who did believe they were fake, most decided the fragments were 19th century reproductions. Not a single student ever suggested that the museum fabricated these artifacts for the purposes of this project. If the museum staff had been trying to create a perfect, flawless cartonnage fragment reproduction, then these numbers might be understandable; indeed, they would testify to the skill of the museum staff members involved. The museum staff, however, had created artifacts specifically designed to be identified as forgeries; clearly visible are several prominently placed clues to the artifacts' lack of authenticity (see below). Amazingly, the students' tendency to find the artifacts authentic was not because they did not find these clues. On the contrary, the students seemed to have an inherent reluctance to pronounce the artifacts fake; they went to tremendous lengths to ignore and/or explain away these clues when they were discovered. This paper will attempt to determine why this occurred.

Methods and Limitations of Study

This is a discussion of preliminary research findings drawn from observations made by museum staff members as they worked with students over the past five years. This is not a formal study with rigorous data collection procedures. When this program began, it was not the intention of the museum staff members to conduct a formal research project. Although extensive data were gathered from the students each time the program was run, identical data collection instruments were not used in each instance.

A significant amount of data, however, is available. The majority of the panel presentations were videotaped, as well as many of the regular classroom sessions where the students interacted with the artifacts and researched ancient Egyptian mummification. The creation of the cartonnage fragment reproduction was carefully documented. Extensive notes were taken throughout the project, by museum staff members as well as teachers. The various versions of the online documents, resource guides, and database records are all well documented. One implementation of the project (which took place during Fall 2000) was thoroughly documented and described by a doctoral student from the College of Education at the University of Illinois as part of a research study into museum outreach programs in public elementary schools (Costantino, 2001).

There are, naturally, limitations to these data. The fact that this project has been conducted at different schools, using different approaches, with different students who have different reactions, with non-standard data collection procedures, means that it is not possible to do a comparative study of why a group of students at one school in one year reached different conclusions from a second group of students at a different school in a different year. The available data, however, have allowed us to conduct our present analysis of a) the history of the cartonnage project as the museum staff evolved and adapted the program over time, and b) the interactions between students, teachers, and museum professionals, and the development of those interactions over time. It is our hope that this analysis will be of use to educators and museum professionals interested in educational outreach projects from museums to schools.

We begin our analysis with a discussion of the evolution of this project over the years from 1997 to 2003. This evolution was an organic, natural process whereby the museum's procedures in implementing the cartonnage project changed as we worked with the teachers and students, adapting each implementation to better suit the needs of the educational community the museum serves.

Historical Evolution (1997 to 2003)

Museum staff members working with the cartonnage project quickly found that they needed to be flexible; they could not present exactly the same project each time the program was run. They had to adapt to new and unpredictable situations, which in turn led to new adaptations and a constantly evolving educational outreach program. This section presents an overview of the primary evolutionary stages of this project, discussing the changes made over time, the rationale behind those changes, and what was learned from making those changes.

As mentioned above, this program has been run ten times at seven different schools since its inception; over 300 students in central Illinois have researched the authenticity of Egyptian cartonnage fragment reproductions. All but one of these programs was run with sixth grade students (the other was implemented with a group of gifted students from grades three through five). Not including the initial prototype version run in December 1997, there have been three distinct versions of this project; the third version is the one that museum staff members are currently implementing in local schools. This section of the paper includes descriptions of each version of the project, and discusses how each version evolved over time in response to the needs of students, teachers, and museum staff members (see Table 1 for a summary of this evolution).

Prototype Version (1997)

In December 1997, museum staff members, working with the sixth-grade teacher whose idea began the project, decided that the most appropriate artifact for this exercise would be a piece of ancient Egyptian cartonnage. Since no actual artifact could be taken to the schools, museum staff members decided to provide online access to the museum's database records for a dozen actual pieces of cartonnage in the Spurlock Museum's collections. The teacher, who still wanted a physical artifact, decided to manufacture a fake piece of cartonnage herself. She took a six-inch square piece of linen, soaked it in plaster, painted it with Egyptian hieroglyphics, and -in an attempt to age it – buried it in her back yard (when delivered to the class, it still had a fresh pine needle stuck in it).

With this prototype version, the vast majority of the work was accomplished by the teacher herself. She developed all the documentation necessary to guide her students' research in problem-based learning. She gathered all the resources her students would need to examine the artifact and determine its authenticity. She incorporated all of this into her ongoing unit on ancient Egypt. Museum staff members made themselves available for email questions, and served at the end of the project on a panel so that the sixth grade students could present the findings of their research. Aside from this, however, the museum played no direct role in managing the prototype version of the project.

Version I (1998-1999)

One year later, this same sixth-grade teacher, along with an additional teacher from a different school, approached museum staff members to see if they could do this project again. This time, museum staff members were able to dedicate a significant amount of time to improving the project; the result was the first full-featured version of the cartonnage project. The major changes from the Prototype to Version I were an improved version of the cartonnage fragment, an entirely new online interface, and a larger role for the museum in the presentation of the problem.

The museum's Collections Manager and her assistant took on the task of creating a new, improved version of the cartonnage reproduction (cf. Figure 1). It took over twenty hours of work for them to design and create a replica mummy pectoral piece using plaster-soaked linen; the pectoral was painted to resemble an Egyptian artifact. To clearly identify this piece as a forgery, they purposely planted several clues on the artifact: spots of hot pink and hot green paint, several non-Egyptian hieroglyphs, and a Mayan figure in full regalia. Once the artifact was complete, it then took another twenty hours of work to destroy the artifact and age it two thousand years. The artifact was divided into six pieces, and two sets of three pieces each were mounted in glass display frames. By creating two sets of fragments, museum staff members made it possible for the program to be run at two different schools at the same time. They ensured that each set included sufficient clues, including half of the Mayan figure. The small, left-over pieces of cartonnage that broke off during the aging process were placed into plastic bags so that students could analyze these fragments in detail.

The museum staff also made major changes to the online interface. Realizing that simple access to database records was not sufficient information for the students, museum staff members created an elaborate, fictional database access system called the "Spurlock Museum Management System." This "system" was simply a series of password-protected ~~eb~~ Web pages which gave the illusion of providing access to detailed information about museum operating procedures. The password provided to the students gave them access to two things: a series of documents about museum careers (including job descriptions and required education levels for each job) and a "main project page" for researching the cartonnage fragments. This project page provided links to four types of resources: a set of research notes, a list of database records, a glossary of terms, and a *WebBoard*. The research notes were a series of lengthy pages about Egyptian mummification practices, linked to multiple pages covering more detailed topics, with a navigation system that challenged students to read the text carefully when searching for answers. The list of database records was an expanded list of relevant Egyptian artifacts from the museum's collections. The *WebBoard* provided a simple discussion board interface where students could post questions and museum educators could post responses.

Spurlock Museum staff members assumed a new, major role in the presentation of the project to the students, making substantial modifications to how the "problem" was described. In this version, students were told that the museum had received a set of Egyptian cartonnage fragments from the registrar at the Sherman County Historical Society. Since this historical society only collected local heritage materials, the registrar wondered if the Spurlock Museum would be interested in accessioning the artifacts, which had been found in an unmarked box in their storeroom. The students were asked to help the museum staff research the artifacts' provenance and determine their authenticity. This setup was, of course, completely fictitious; there is no "Sherman County Historical Society." Spurlock Museum educators created a letter from the fictitious donor (see Appendix 1) who supposedly gave the artifacts to the historical society and a letter from the Spurlock Museum's director, thanking the students for their help in this research project. Museum staff members also created an accession card for the artifacts, listing the information supposedly recorded by the historical society. The letter from the donor was specifically written to address the UNESCO treaty issue, and provided evidence that the artifacts entered the United States before 1966. Copies of these documents were delivered to the students participating in the program.

Museum staff members also assumed a major role in delivering the artifacts to the students. The cartonnage fragments arrived with a lot of "pomp and circumstance," carried in by the museum's Collections Manager and her assistant wearing white lab coats and using gloves. The museum educators even considered having a security guard accompany the artifacts in full uniform, but decided against it.

Version 2 (2000-2001)

Early in the year 2000, museum staff members decided that they needed to make additional revisions in the cartonnage project. Although the physical artifacts were a great success, the online resources were being insufficiently used and needed to be improved. Moreover, museum educators were unsatisfied with the presentation of the problem to the students, as the fictitious setup was making it difficult to run the same project for a second year in the same school. Attempts to resolve these issues resulted in Version 2 of the cartonnage project.

For Version 2, the online resources available to the students were completely revised. The fictional "Spurlock Museum Management System" was dismantled and replaced with a simple Web page that linked to three resources. The first link was to a resource guide about Egyptian Mummification (http://www.spurlock.uiuc.edu/education/resources/mummification). The second link was to a series of documents about the research process (http://www.spurlock.uiuc.edu/education/resources). The third link was to the same *WebBoard* from Version 1 (no changes were made to the *WebBoard* at this time).

The original version of the resource guide had proved so difficult to navigate that it was unusable by most of the students. The online resource guide for Version 2 was completely reorganized, with a new graphical interface that included illustrations and embedded links for further information about relevant artifacts in the museum's collections. The guide was organized into seven categories: history, rituals, artifacts, sources, materials, chronology, and glossary (see Figure 2 for a sample screen shot). An external graphic designer was hired by the museum to develop the site's look and feel, and construct icons and other images.

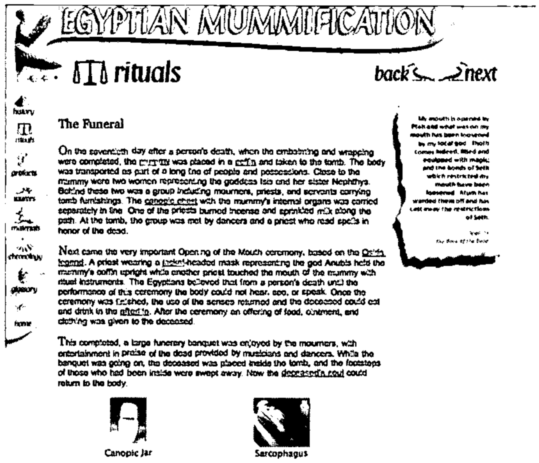

Figure 2: Sample Screen from Online Resource Guide

In addition, museum educators wrote a series of documents about the research process (including information on identifying reliable sources and distinguishing fakes from forgeries); these documents were placed online in Adobe Acrobat format. These additional resources were specifically requested by the teachers, who wanted to provide their students with more accurate information about how to conduct research and how to identify fakes and forgeries. In writing these documents, museum educators took advantage of this request to remedy existing problems with the program. One document, for example, helped the students refocus their research efforts (which often strayed from the artifacts) by teaching them how to observe the artifacts more carefully. A second document addressed the common student belief that if something looked real, then it was; this document explained how forgers were skilled at making things look real, even if they were not. Finally, to help the students understand the research process even better, museum educators developed a "research packet," which helped the students formulate their research questions, record their findings, and document their sources (see Appendix 2).

With the new online resources available, museum educators attempted to sort out the difficulties behind presenting a fictitious problem to a class of students. What resulted was a mishmash of techniques, where over the next several implementations of the program museum staff presented the problem in slightly different ways. They still claimed that the object was from another museum, but presented it as an educational program specifically designed to allow school children to work with museum professionals on real life projects. Other than the letter from the donor, no additional documents were given to the students. Within the donor letter itself, any references to the name of the artifacts' original owner or the name of the "Sherman County Historical Society" were blacked out. The students were told that the letter was censored to protect the privacy of the donor; in reality, the censoring helped museum educators avoid a growing problem where students kept trying to locate the "Sherman County Historical Society" in real life.

Version 3 (2002-2003)

By 2002, the museum educators – although still happy with the physical artifacts and online resources – continued to have difficulties with the presentation of the problem to the students. Since this project was now run every year, often at the same school, museum educators wondered if the students would soon stop believing that the museum still needed their help researching the same artifact. To resolve this dilemma, museum staff members decided to return to the original idea of the sixth-grade teacher who initiated the project in 1997.

Version 3, the current version of the project, presents the situation as a problem that happens to curators all the time. Museum staff members tell the students that while the cartonnage fragments were on display in their museum, a stranger claimed that the label copy describing the artifacts was erroneous. The students are then told that this stranger left immediately and that the museum had no way of determining his identity or asking him any follow-up questions. The students are then asked to help the museum staff members confirm the contents of the label copy (see Appendix 3). By presenting the problem in this way, the museum educators have found it much easier to deal with direct questions from the students with much less fabrication. If the students ask the direct question "are you telling us the truth?" the museum educators now reply that the setup (that a stranger commented on these artifacts) is not true, but that these are actual museum artifacts and that this is an actual problem which curators deal with on a regular basis.

Finally, two minor changes were made to the online resources during the transition from Version 2 to Version 3. First, museum staff members decided to get rid of the simple, introductory page linking to the online resources; all online materials are now accessible from the Programs and Events section of the main Spurlock Museum Web site (http://www.spurlock.uiuc.edu/education). Second, museum staff members decided to no longer use the *WebBoard* for communication with the students. Although this may seem like a major change to the online interface, the *WebBoard* discussions were quite easily replaced by email (see below for further discussion).

This discussion has therefore summarized the evolution of the different versions of the cartonnage project. One implementation was run using the prototype version. Three implementations were run using Version 1. Four implementations were run

	Dates	Artifacts Made by	Documentation	Online Resources	Communication
Prototype	1997	Teacher	No formal documents	Stand alone database records	Email, and panel presentations
Version 1	1998-1999	Museum	Accession card, letter from fictitious donor, and letter from museum director	Fictional museum management system with integrated database records and early version of resource guide	Staff visits, email, webboard, and panel presentations
Version 2	2000-2001	Museum	Letter from fictitious donor (censored)	Complete resource guide with integrated database records as well as research documents	Staff visits, email, webboard, and panel presentations
Version 3	2002-2003	Museum	Letter from fictitious donor (censored); exhibit label copy	Complete resource guide with integrated database records as well as research documents	Staff visits, email, and panel presentations

Table 1: Summary of Evolution of Versions

using Version 2. Two implementations were run using Version 3. At least two more implementations using this version are scheduled for the spring semester of 2003.

Analysis of Student Activities and Interactions

This section presents an analysis of the students' activities and interactions as they worked with the artifacts, the available resources (both physical and online), their teachers, the museum personnel, and each other. It discusses the overall trends we observed in terms of how the students participated with the museum personnel on the project, what their levels of involvement and understanding of the project were, and how these things changed over time. In doing so, we will attempt to determine what the students learned from participating in this project.

There will be four different types of activities and interactions discussed. First, we will look at how the students interacted with the artifacts. Second, we will explore how the students made use of the information resources, both physical and virtual, at their disposal. Third, we will examine the different approaches to research taken by the students. Finally, we will discuss the students' learning processes, focusing on how their levels of understanding changed over time.

Interactions with Artifacts

How did the students interact with the artifacts? The artifacts were generally brought to the school and presented to the students fairly early in the project; the students would usually have had some initial introduction to ancient Egypt, have been told about the upcoming project, and have been looking forward to the arrival of the actual artifacts themselves. The artifacts were brought to the school by museum staff members who clearly treated the artifacts with a great deal of respect. The teachers were required to sign an official "loan form," further impressing upon the students the value of these artifacts. The students were taught proper artifact handling procedures which stressed the importance of handling the artifacts carefully, with gloves, and treating the artifacts with respect at all times. The students were not allowed to handle the main pieces of cartonnage; these were kept intact in a glass display frame. The students were allowed to handle several small pieces of cartonnage, but only with tweezers. The seriousness of these procedures had a big impact on the students. They were very careful to always wear gloves when handling the artifacts; they treated the artifacts with respect, as they had been trained to do. During the panel presentations, if some students forgot to wear gloves, other students would even remind them to do so. The extent of their training in artifact handling likely was a major factor in the students' persistent belief that the artifacts were authentic; if the artifacts were not authentic, some students asked us, why would we be so strict about artifact handling procedures? The students never asked why, however, if the artifacts were authentic, museum staff members would have let them out of the museum.

What did the students do while studying the artifacts? As will be discussed in more detail below, students took a common approach to studying the artifacts, examining the hieroglyphics, colors, gods, and material composition of the pieces. They tried to identify the gods, decipher the hieroglyphics, and determine whether or not the colors on the cartonnage fragments matched known pigments available to the ancient Egyptians. Since students were unable to remove the cartonnage fragments from the display cases, they had to rely on the fragments provided in the plastic bags for material analysis. They typically examined these small fragments carefully to see what they were made of and determine their authenticity and, if possible, the age of the fragments. Only in a few cases did students attempt anything innovative; one group, for example, extrapolated from the fragments they were given a theoretical reproduction of what the entire pectoral piece might have looked like.

How much time did the students spend working directly with the artifacts? In most cases, students did not spend as much time working directly with the artifacts as the museum staff members expected they would. Their typical approach was to spend one day studying the cartonnage fragments at the beginning of the project, and then put the fragments aside and only occasionally look at them before the final panel presentation. The students spent enough time with the objects to answer questions about the appropriateness of the colors and attempt a translation of the hieroglyphics, but the artifacts themselves were rarely the focus of the classes' attention. After their initial study of the artifacts, the students typically turned their attention to the research process, searching the Internet and their school's library for information about cartonnage and the mummification process. Only recently have museum staff members made a point of stressing in class the value of constantly returning to the artifacts during the research process. While having the artifacts in the classroom throughout the project provided a catalyst for the students' research activities, the students did not understand the significance of making the artifacts a part of their research process until they were told to do so.

Use of Information Resources

How did the students make use of the online resources provided by the museum? Depending on computer availability, students were regular visitors to the museum's Web site and resource guide; almost every student referred to the museum's Web site in their panel presentations. It appears as if the fact that museum staff members provided them with an online resource guide to Egyptian Mummification had a significant impact on the students' approach to the project. The online resource guide helped them start the research process, and they searched the museum's Web site for clues as to the artifact's authenticity. Many students believed that the mere fact that the museum had provided this Web site proved that the artifacts were from ancient Egypt. Why else would we have given them that site? Students stopped asking this question when museum staff members decided to make these online resources available from the museum's public Web site along with resource pages on several other topics. Nevertheless, as the source of both the artifacts and a primary source of data about Egyptian mummification, the students usually accepted the authority of the museum without question.

What other resources did the students use? All students had access to a variety of information resources, including books, articles, and Web sites, accessible from their classroom, their school or public library, or their home. The students were frequent users of the open Internet; in the late 1990s, some students relied almost entirely on the Internet for the data. Every student used at least one Web site (other than the site provided by the museum) while conducting the research. Most often, these were sites on the use of color or hieroglyphics in ancient Egypt. Different classes almost always located completely different Web sites, which were passed from one group of students to another, and on which students often based the foundations of their arguments. In this way, Web sites with erroneous information frequently spread mistakes throughout an entire classroom of students like some kind of knowledge virus. Students also used books (although only infrequently would students use their textbooks) in their research, often looking up translations of hieroglyphics or trying to identify the images on the fragments. Many students also found basic information on ancient Egypt and mummification from books; in one case, a group of students working on this project checked out almost every book on ancient Egypt from their local public library.

Approach to Research

What research questions did the students ask? On the first day of the project, students were encouraged to develop a list of questions they would like answered. More questions would generally be added to the list later on as they examined the artifacts and spent more time conducting research. The most commonly asked questions were: Who made this? When was it made? What was it made for? What is it made from? What time period is it from? Where is it from? What is shown in the pictures? What do the hieroglyphs say? What is the significance of the colors? Over the course of the project, however, the students usually became more and more focused on the overall question: Is this real?

What research methods did the students use to answer these questions? To ensure that the students started by looking at a source they could count on as reliable, museum staff members initiated the research process by providing the address of the Museum's online resource guide to Egyptian Mummification. Most often, students began the project by spending time either in a computer lab or on one or more computers in the classroom. The more available computers, the more likely students were to focus on online resources. Having fewer available computers quickly directed the students toward their library books. Museum staff members were available for questions over the Internet (via email or on the *WebBoard*). In addition, at least one museum staff member would visit the students for a question/answer session halfway through the project. To provide the students with a consistent source from the museum, the assistant director of education was established as the students' primary contact at the museum.

What approach did the students use to determine the authenticity of the cartonnage fragments? With few exceptions, students primarily focused on four factors: the hieroglyphics written on the fragments, the colors painted on the fragments, the materials the fragments were made from, and the gods depicted on the fragments. These factors, almost always in this order, comprised the four main arguments put forth by the students in the panel presentations. Based on their findings regarding the hieroglyphics, colors, material types, and the gods, the students would render a verdict of authentic or forgery in front of a panel of museum staff members.

What was the most frequent decision reached by the students researching the fragments? As discussed above, about three out of every four students decided that the artifacts were authentic; of these students, about half reached this conclusion based solely on their analysis of the colors found on the artifacts. Most students were absolutely convinced that these reproductions were real, and this belief was not limited to the reproduction made by the museum staff. When the prototype version was run, one student told us that he believed the museum staff had taken a piece of cartonnage from their collection and ripped a square piece out of it; he even showed us, on the online database, a piece of cartonnage that – by complete coincidence – was missing a piece exactly the right size and shape. The major obstacle the students faced throughout the project was their tendency to look for reasons to prove that these artifacts were real, and not to prove that they were fake. Even when faced with the clues planted by museum staff (such as the spot of hot pink paint), students tended to rationalize these clues away (arguing that perhaps red pigments, when two thousand years old, might look like hot pink).

Did any student ever discover the Mayan figure? From 1997 to 2003, the Mayan figure has only been discovered three times, and each time it was discarded and never brought up again. The first time this happened, the student who discovered it pointed out that this looked like a Mayan figure but decided that this was not a problem since it was likely that Egyptian and Mayan cultures were aware of each other; she told us that she reached this conclusion because she had learned about Egyptian and Mayan civilizations from the same book. The second time the figure was discovered, it was identified by a learning disabled student who thought there was something odd about that figure, but could not decide exactly what it was; before he could explore it further, however, he was quickly shut down by his fellow group members. The third and final time the Mayan figure was discovered, a group of students thought the figure looked like an Aztec warrior; they explained away this apparent anomaly, however, by stating that since the Egyptians and the Aztecs both lived in the tropics, it was likely that the Egyptians also portrayed figures with leaves around their wrists and ankles, as the Mayan figure was depicted.

Learning Processes

How did the students' understanding of museum activities change over time? The students' understanding of what a museum curator does improved dramatically over the evolution of this project. In Version 1 of the program, the students had a very hard time understanding the role of the museum curator when working with artifacts. This was, in part, because the museum educators, in an attempt to provide a certain amount of intellectual rigor, originally told the students that museums were expected to possess 100% accurate information, and that museum curators were responsible for ensuring the accuracy of the museum's information. Students were provided a list of "questions curators were supposed to be able to answer," including questions about the physical characteristics of the artifacts, the provenance of the artifacts, whether or not the artifacts complied with the 1971 UNESCO treaty, and the basis for information about the artifacts as provided by the donor. This list was supposed to help the students formulate their own research questions; instead it became a burden, forcing students to worry about treaties instead of the artifacts. To solve this problem, the museum educators adapted their approach in Version 2, asking the students to develop their own questions through small group discussions. Encouraging the students to think like detectives helped them become inspired by the job of the museum curator. The new presentation of the problem developed in Version 3 went even further, providing the students with an even more straightforward problem to solve: who do you think is right? The mysterious stranger? Or the donor? By drawing the students into the inquiry process, these changes helped the students better understand the job of the museum curator.

How did students' understanding of the research process change over time? The students' abilities to conduct research and understand research methods improved dramatically from 1997 to 2003. This was especially true in the case of the students' online searching skills, where they became much more aware of the need to assess the quality of information they found online. In December of 1997, one student actually believed (and stated) that everything she read on the computer had to be true. By the spring of 2002, however, students were comparing the overall reliability of sites from dot-edu or dot-com domains. This change is likely due in part to the efforts of both museum professionals and teachers to explain the importance of good research methods to the students. Throughout the development of this program, for instance, museum staff members have continued to provide more explicit documents about good research habits. Museum educators and teachers alike continued to stress the importance of having and citing a source for every piece of information. These efforts have consistently led to students with higher quality research skills.

How did the students' approach to logical arguments and critical thinking change over time? The students' abilities to think critically and make logical arguments did not change significantly over the course of this study. Students did not want to believe that the museum might bring them artifacts that were not real, and they would go to tremendous lengths to dismiss any evidence that might possibly imply a fake. When arguing that the artifacts were authentic, students twisted their logic in strange, clever, and often unexpected ways; students who thought the artifacts were fake, however, usually did not employ such logical contortions. For instance, when attempting to translate the hieroglyphics on the artifacts, students usually found several hieroglyphs they could not identify: the signs planted by the museum staff as clues that the objects were fake (some of the "Egyptian hieroglyphs" were actually Mayan numbers). Instead of reaching that conclusion, however, students decided that these hieroglyphs must either be signs that had not yet been discovered or signs that were too rare to be listed in the books in their school library ("we're only sixth graders, after all," one group told us). The possibility that these signs might not be Egyptian hieroglyphs at all never occurred to them. For another example, when attempting to use the colors found on the artifacts to determine their authenticity, the students almost always began with a list of colors the ancient Egyptians could produce and then looked to see if those colors were found on the artifacts. They almost never checked the reverse, looking to see if any colors found on the artifacts could not have been made in ancient Egypt. In ways such as these, the students would skew their answers to match their expectations.

Analysis of the Museum/School Relationship

This section presents an analysis of what we learned about supporting a cooperative relationship between museums and schools when conducting educational outreach programs. It discusses what we learned as we moved from version to version, evolving the program to create a better product. In doing so, we hope to summarize the most important lessons we learned about what makes for a successful relationship between museum educators and school teachers.

We will examine three different areas of this relationship. First, we will look at levels of involvement in the classroom, from the perspectives of both teacher and museum educator. Second, we will examine the issue of availability of online resources and the importance of access to technology. Finally, we will discuss the relevance of this program to the curricular needs of the school and its impact on the success of the project.

Levels of Involvement in Classroom

How involved were museum staff members in classroom activities? The success of this program was directly related to the amount of time museum staff members were able to dedicate to the classes researching the cartonnage fragments. From the onset of Version 1, Spurlock Museum educators and other staff members remained very involved in the classroom: they made several trips to each class, at the beginning, middle, and end of the project; they were constantly available to answer questions for the students, either by email or using the *WebBoard*. Any museum educator planning to initiate a project such as this one must be aware of the need to dedicate a great deal of time to helping the students have a successful experience. In addition, when collaborating with teachers who were not so familiar with active learning or problem-based learning techniques, the museum educators found that they had to spend more time in the classroom and play more significant roles in setting up the project and assisting the students.

How involved were the teachers in classroom activities? The success of the program was directly related to the amount of time the teachers were able to spend with their students on this project. The more involved the teachers were, the more confident they were in their research skills, and the greater their knowledge of the content area, the more the students got out of the project – in terms of Egyptian history, museum practices, and research activities. Likewise, the more time the teacher spent researching ancient Egypt ahead of time, finding appropriate resources and checking books out of the library in advance, the greater the impact of the project on the students. Students who were able to spend three weeks on the project generally learned more than students who only had two weeks; this was especially true with schools taking a problem-based learning approach (which took longer to set up). Few schools, however, could afford to dedicate three entire weeks to this project. Nevertheless, even teachers who could spend only one week on this project found it instrumental in changing the perception of the museum from a once-a-year field trip destination to a valuable resource that could be directly integrated into the classroom curriculum. With the students relying heavily on the Internet for their information, however, teachers had to be prepared to vet for inappropriate material all sites discovered by the students.

Availability of Online Resources

How did access to the online resources provided by the museum affect the students' successful completion of the project? Students working with the cartonnage fragments had access to museum-provided information about Egyptian mummification practices (including information specifically about cartonnage) as well as museum practices and research methods. When students did not have access to this information, as in the prototype version when only an online list of database records was provided, students gathered most of their information from other online sources. As the museum made more information available online, and as artifact database records were integrated into the online resource guide, the students began to rely on the museum-provided Web site as their primary source of information. Students were not particularly interested in the elaborate, fictional "Spurlock Museum Management System" provided in Version 1; they preferred the straightforward access to information about mummification and research activities available in Version 2. Students, however, were much more likely to make use of the resource guide on Egyptian mummification than they were to read the documents written by the museum educators about conducting research and determining the authenticity of an artifact. It is possible that this indicates a preference for information provided in a graphical format over text-only documents. It is also possible that this reflects the students' (or teachers') unfamiliarity with Adobe Acrobat; to account for this possibility, the museum educators have recently begun taking hard copies of the research documents with them on their classroom visits.

How did access to the necessary technology affect the program? Although successful completion of the project assumed that the students would have access to the Internet in order to read the online resources provided by the museum, this was not explicitly a requirement of the program. Several schools had to cope with technological failures that resulted in students having little or no time in the computer lab. Students who were unable to access the museum's online resources had to rely on their school libraries for information. When the museum was working with students who had access to the Internet in their classroom, the museum staff found that the students were more likely to have a thorough understanding of what the museum wanted them to do (especially when working computers were available throughout the project). These students, however, were also more likely to have found less reliable sources of information by concentrating their search on the open Internet.

What was the impact of using the *WebBoard* for communication between students and museum educators? With the implementation of Version 1, museum staff members had very high hopes for the museum-school interactions made possible by the *WebBoard*. At the end of Version 2, however, the museum staff decided no longer to use the *WebBoard* technology. This was for several reasons. First, Spurlock staff members were not running the *WebBoard* from their own server, and therefore were paying expensive licensing fees each time the program was run. Second, the *WebBoard* discussions often turned out to be unhelpful and sometimes harmful to the project. Students wasted a great deal of time making off topic or inappropriate posts. They would usually ask the same questions over and over, wasting the museum educators' time. They would post useless requests such as, "Please tell us whether this is authentic. Thank you." Some students would even post threats, demanding that museum staff members tell them the answers "or else." In the end, it was simply easier for the museum to drop the *WebBoard* component and rely solely on email for asynchronous interactions between students and museum employees. To reduce the number of inappropriate or repetitious questions, the teacher now collects the students' questions, are now collected by the teacher, who edits them, and sends a moderated list to the museum educators. This illustrates the advantage of having greater teacher involvement in online communication between schools and museums.

Relevance to Curriculum

How did this lesson relate to classroom projects as part of a course on ancient Egypt? One of the biggest factors behind the success of this project was that it is most often used to introduce a unit on ancient Egypt. Providing the students with a hands-on, active learning experience made the concept of studying Egypt more exciting. Having a practical introduction such as the cartonnage project made the students look forward to learning more about Egypt in the traditional classroom setting.

What was the value of the panel presentations to the project? The fact that the students were required to do group presentations to a panel of museum staff members at the end of each project had a serious impact on the students' behavior. Knowing that the museum staff would be "evaluating their work" made them take the project very seriously. In addition, the panel discussion format let the students know that the museum took their work seriously as well. The museum educators explained to the students that no one expected them to provide all the answers to every question. Nevertheless, the fact that the museum employees actually listened to the students and took their advice to heart clearly impressed the students and made them feel that their work was truly important.

Is this project beyond the capabilities of elementary or middle school students? Over the years, the museum staff has adjusted and adapted the program to reduce the number of potential pitfalls that could trip up the students. Most recently, museum educators dropped the discussion of UNESCO treaty issues simply because it was too hard for most sixth grade students to understand. The one time that the museum tried running this project with elementary school students (gifted students from grades three to five), the students had trouble understanding the purpose of research; they were also plagued by technological problems which limited their time on the computers. Even at that young age, however, and despite challenges of time and technology, the project successfully motivated the students and piqued their interest in both ancient Egypt and research methods. For middle school students, the program, at least in its current form, seems to be at an appropriate level for an introductory project to ancient Egypt, the museum profession, and the practice of conducting research.

Conclusions

The evolution of the cartonnage project from the prototype version in 1997 to the current version in 2003 was extremely time-consuming for the staff of the Spurlock Museum, requiring a great deal of analysis and revision. They worked toward a program that would be easy for teachers to manage and not too hard for students to understand. They worked to help students avoid pitfalls and come to the right conclusions. Their most difficult challenge, however, was getting the students to accept the project on their own terms and at face value.

It was only recently that students were willing to accept the project as a simulation, to approach it as a learning experience. In doing so, the students were freed to understand the true goal of the exercise: to appreciate the purpose of research. They were no longer looking for the right answer to fill in a blank on a test; rather they were coming to their own conclusions. They understood that there was no trick question, that the right answer was the answer they uncovered as researchers, that they themselves were the source of the answers. Moreover, they understood that the answer itself is not as important as the process of inquiry: asking questions, conducting research, and supporting their findings with sources.

The amount of time required to reach this goal demonstrates the difficulty of creating a program such as this; museum staff members needed to keep evolving, improving, and constantly working to meet the needs of students and teachers. Museum educators had to wrestle with difficult questions. In using fictitious documents, did they risk their own credibility while emphasizing the importance of credible and reliable sources? Did they have an obligation to let the students know eventually that these artifacts were in fact reproductions? One group of students did track the museum educators down, months after the end of their project, demanding to know the answer... and they were told. But those students were the only ones that ever reacted in that way, the only group that was overly bothered by the fact that museum educators were perhaps not telling them the whole story. And this, for the museum staff, best illustrates the success of the cartonnage project. When students see that they themselves are the source of the answers, that there is no right or wrong, that there is only a process of discovery, students learn the only real truth.

References

Adams, C., T. Cole, C. DePaolo, & S. Edwards(2001). Bringing the curatorial process to the Web. In D. Bearman & J. Trant (Eds.), *Museums and the Web 2001* (pp. 11-22). Pittsburgh, PA: Archives and Museum Informatics.

Bennet, N. & B. Trofanenko (2002). Digital primary source materials in the classroom. In D. Bearman & J. Trant (Eds.), *Museums and the Web 2002* (pp. 149-156). Pittsburgh, PA: Archives and Museum Informatics.

Bennet, N., & B. Sandore (2001). The IMLS digital cultural heritage community project: A case study of tools for effective project management and collaboration. *First Monday*, 6 (7).

Boud, D. & G. Feletti, (Eds.) (1991). *The challenge of problem-based learning*. New York: St. Martin's Press.

Costantino, T. (2001). *The Museum Problems in Today's World Program – Research: Ancient Egypt Module*. Research paper written for the College of Education, University of Illinois at Urbana-Champaign.

Frost, C. O. (1999). Cultural heritage outreach and museum/school partnerships: Initiatives at the School of Information, University of Michigan. In D. Bearman & J. Trant (Eds.), *Museums and the Web 1999* (pp. 223-229). Pittsburgh, PA: Archives and Museum Informatics.

Frost, C. O. (2001). Engaging museums, content specialists, educators, and information specialists: a model and examples. In D. Bearman & J. Trant (Eds.), *Museums and the Web 2001* (pp. 177-188). Pittsburgh, PA: Archives and Museum Informatics.

Marty, P. (2000). Online exhibit design: the sociotechnological impact of building a museum over the world wide web. *Journal of the American Society for Information Science*, 51 (1), 24-32.

Sumption, K. (2001). "Beyond museum walls" — a critical analysis of emerging approaches to museum web-based education. In D. Bearman & J. Trant (Eds.), *Museums and the ~~Web~~ Web 2001* (pp. 155-162). Pittsburgh, PA: Archives and Museum Informatics.

Teather, L., & K. Wilhelm (1999). Web musing: Evaluating museums on the Web from learning theory to museology. In D. Bearman & J. Trant (Eds.), *Museums and the Web 1999* (pp. 131-143). Pittsburgh, PA: Archives and Museum Informatics.

Appendix A: Donor Letter (redacted)

[redacted]

[redacted]

June 21, 1966

Dear [redacted]:

Thank you for taking the time to talk to me on the telephone yesterday. I hope the Egyptian objects I am sending you will be of use to the [redacted].

You asked me to give you all the information I have on these objects for your records. It isn't very much, I'm afraid. They are an inheritance from my great-uncle, who lived in Scotland, and my only visits with him took place many years ago. In truth, I am very surprised and pleased to find he had remembered my childhood fascination with his collections after so many years. Here is what I remember of what he told me:

1) The objects were purchased by my great-uncle, [redacted], in Egypt during the 1890's. I remember him saying they had "cost him a pretty penny," but that he never regretted the cost.

2) The man he bought them from told him that the pieces were from a mummy's painted chest decoration. The man also told him that, in life, the person had been a high priest and advisor to the pharaoh.

3) My great-uncle had always talked about taking the objects to a museum to have an expert date them, but to my knowledge he never did.

I hope this information is of some help to you. If you have any other questions, please let me know. I will be happy to help if I can.

Sincerely,

[redacted]

BEST COPY AVAILABLE

Appendix B: Research Packet (sample page)

Research Notes
Spurlock Museum

Beginning Date: ________________ **Report due:** ____________________
Researcher(s): ____

Note: Every hypothesis must be supported by references from reliable sources. All source information must be recorded in detail (i.e., bibliographic entry and page numbers).

Preliminary Artifact Examination Notes and Questions:

Research Areas:
1)

Resources: ____

Appendix C: Exhibit Label Copy

Bound for Eternity

These fragments are thought to be from an ancient Egyptian mummy's painted chest decoration.

Fragments.
Egypt, unknown date. Cloth, plaster, pigment.
Anonymous donor. 1966.3.4-6

Focus your young visitors: -Kids Innovation- Fundamental changes in digital edutainment

Sebastian Sauer, ion2s buero fuer interaktion, and Stefan Göbel, ZGDV e.V. Digital Storytelling, Germany

Abstract

With regard to the acceptance of human-computer interfaces, immersion represents one of the most important methods for attracting young visitors into museum exhibitions. Exciting and diversely presented content as well as intuitive, natural and human-like interfaces are indispensable to bind users to an interactive system with real and digital parts.In order to overcome the obstacles to attracting the young, we have taken an interdisciplinary approach to the fields of multimedia, media design, TV/movie or literature, and especially myths, to develop narrative edutainment applications for kids. Similar to Hollywood, production content and dramaturgic scripts are generated by authoring tools and transmitted by multimodal interfaces enhanced by natural conversation forms.These concepts are enhanced by interaction design principles, methods and appliances; such as Kids Innovation or Momuna.This paper describes both existing technologies and new innovative methods and concepts of interactive digital storytelling and user interaction design to establish immersive edutainment applications and appliances using multimodal user interfaces.These approaches merging into Dino-Hunter as innovative game-oriented mobile edutainment appliances are discussed within the environment of mobile application scenarios for museums and their young visitors.

Keywords: Kids Innovation, Momuna, Interactive Digital Storytelling, Dino-Hunter, Telebuddy, GEIST, Experience Appliances, Transforming User Interface, Multimodal Interfaces, Digital Information Booth, Edutainment Applications.

Introduction

The global aims and tasks of museums are not limited to the archival storage and conservation of cultural and scientific exhibits and artifacts. Rather, the transmission of knowledge, the imprint of cultural identity and the experience of art represent special issues to museums and involved parties in the museum scenario. In order to be an active part of cultural life, it is necessary to compete with other actual offers and trend-setting developments in the current information society and to influence these developments. This paper describes practical examples, concepts and methods of innovative digital and interactive edutainment applications enabling especially young visitors to dive into a game-oriented experience and learning environment within museums. Examples for this are integrated multimedia systems of experience appliances consisting of avatar-based information kiosks, mixed reality environments, interactive toy characters and mobile information devices, all based on alternative trend-setting user interfaces providing multimodal and adventure-driven access to the contents of museums. The common scenario and introduced new concept of Kids Innovation is modularly structured and addresses exactly these issues mentioned above. Altogether, Kids Innovation focuses on young people and provides mechanisms and methodologies for fundamental changes in digital edutainment applications. Thus, new ways of interaction and communication between kids and artifacts, between visitors and the museum, or between kids and visitors are introduced.

From the technical point of view, Kids Innovation is enhanced by different principles and mechanisms of interactive digital storytelling and user-centered interaction design methods. Interactive storytelling techniques enable museum institutions to provide exciting stories and interesting presentation modes "beyond the desktop", enabling (young) visitors to dive into a story and get detailed information about artifacts using game-based interfaces. These game-based interfaces are integrated within a mixed reality environment consisting of physical interfaces/ props and devices such as sensors, video-tracking, speech recognition or scanners, as well as virtual components such as a 3D-environment with narrative characters talking to the visitors, and multimedia presentations. Subsequently, interactive storytelling techniques improve the usability of digital edutainment applications and enable the creation of exiting, suspenseful and immersive interfaces.

From the commercial and marketing-oriented point of view, the use of experience appliances developed by the principles of Digital Storytelling and of Kids Innovation enables museum institutions to get valuable feedback concerning the impact and effect of their own exhibition concepts and the behavior of (young) visitors. Based on these facts and additional case studies, museum institutions are enabled to support schools, classes or any other kind of visitor group within the three major phases of visiting a museum: preparation, execution and post-processing.

State-of-the Art

User Centred Interaction Design

New high-end technology in the field of computers, multimedia systems and telecommunication arises every day and promises a better and easier way to manage human beings' lives. The advantages of such technologies are their nearly never-ending funcitionalities and possibilities. But the crux is, that their complexity is mainly in contrast to intuitive access of their functions. The human being has needs and wants to solve problems without being stressed by having to understand extensive technologies. To let the human being become the centre of the considerations during the development process of technological innovations is the goal of User Centred Interaction Design. Scientists, product and communication designers, computer engineers, psychologists and anthropologists combine their competences to make things usable.

Recently, most multimedia systems exclusively have worked with the combination of haptic hardware (solid-) and traditional graphical user interfaces. You find such user interaction components and principles combined in standard Desktop PCs, Personal Digital Assistants, TabletPCs and thousands of software applications. Sometimes auditory components are integrated, but in a very reduced way. A combined solid-/audio- or audio-only interface is associated with devices which are normally developed for physical handicapped people (for example, the blind) or which are used in a special context, where visual attention cannot be paid to the device; for example, while driving a car.

Nowadays users get off of their office desktop and dive into a mobile situation. That fact is reflected in new industry interactive and digital products for mobile usage, such as smartphones, Pocket PCs or CarPCs. This is an interesting and important point for the User Interface Developer. The mobile situation and the permanent changes of user surroundings cannot be clearly defined as in a more constant environment; for example, at the desktop in the office. The knowledge of how the five senses of the human being work together and how their reactions to different impressions influence the interaction between the user and a device is more important then ever. This point is very interesting within the museum context because of the combination of real artefacts in interaction with the digital appliances.

User Interface combinations and principles show that an assimilation of information can be done parallel by our senses. But every perception channel can support only one orientation of our attention. Jef Raskin [13] picked up on this problem and introduces the "locus of attention" in combination with the term "singularity" [10]. If the User is involved in one linear thinking process, can an unforeseen event deflect attention in a new way? The decision is that no second locus of attention arises, but the old one breaks down and the new one takes priority.

This fact is very interesting for User Interaction Specialists. It influences especially the design process of human device interfaces for mobile multimedia systems because of the constantly changing environment of the user in the mobile situation. The conclusion is that multimodal Interfaces can provide different access to information and allow different communication modes, but you have to understand the temporary context of the user, his physical and emotional possibilities, the variable environment and the dynamic situation. That means that a static multimodal user interface which interacts with a user by constant modes ignores the changes and cannot be so effective as it could be. (For more on this, see "new concepts...").

Digital Storytelling

From the content-related point of view, Digital Storytelling represents a new research discipline contributing to information technology by applying aspects of storytelling and entertainment to digital applications. New authoring methods for multimedia design and application development are currently being explored and implemented by an interdisciplinary team of communication designers and computer scientists. The goal of digital storytelling is to achieve new types of applications that are easily understood by everyone through telling a story [16]. Here, additionally, those new forms of multimodal interaction are integrated. For integrated authoring, three aspects of interactive multimedia play roles in a scenario design description:

- Graphic-interactive media: These are 2D or 3D graphical elements (including spatial sound sources) which can be placed, designed, and depicted in various displays, as well as manipulated by navigation and pointing techniques (an area already represented by Screen Designers and 3D Designers for virtual reality).

- Continous media: These include all kinds of time dependent media, such as audio/video streams, animation, and speech. Design criteria are not questions of spatial layout, but of sequencing and timing as well as turn-taking in a conversation. They determine the dramaturgy of an application to a high degree.

- Physical Media & Modalities: Selection and design of physical-spatial worlds of adequate input and output devices, considering new requirements such as mobility and ubiquitous computing beyond pure screen design. Ergonomic and haptic qualities determine usability, acceptance, and aesthetics.

Designers must integrate, design, and evaluate all of these elements during the concept phase of a production. Ultimately, it is the end-user who will evaluate the overall impression of these integrated aspects and determine the application's usability and value.

To demonstrate the potential of this integrated design, ZGDV Darmstadt has developed a kiosk system for Siggraph 1999 that gives visitors the impression of stopping at a virtual trade show booth where they can get information by having a conversation with a human-like assistant in the form of an avatar [2], a more or less "intelligent" being, a "simple digital assistant"; or by participating with "real intelligence with knowledge and emotions" as an autonomous virtual character [14][8]. This system includes natural interaction elements such as tracking visitors' poses and positions, and the positions of physical objects (e.g. real books), to establish a context for the conversation.

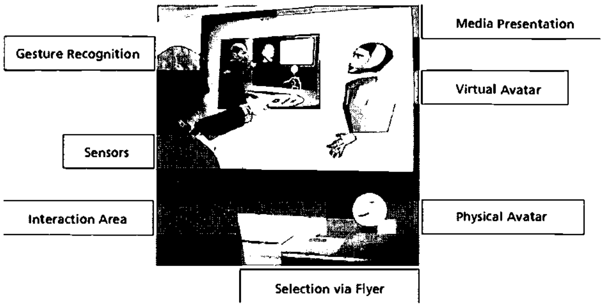

Figure 1: Digital Information Booth – Project "Info zum Anfassen" (tangible information).

Based on that scenario, further innovative ideas and concepts are indicated by the IZA ("Info zum Anfassen", engl.: tangible information) project[1], which aims to develop a mixed reality platform for a booth scenario in cooperation with the visual computing department. Visitors are recognized by sensors and are saluted at the booth by avatars. Corresponding to the application scenario and conversation objectives, a talk is initiated. Here, instead of typical hypermedia, mouse- or menu-oriented interfaces, users can activate sensors or can physically select flyers or can show something, somewhere (gesture recognition) in order to select presentation modes or control the workflow of the story. The story engine is generic in the sense of different application scenarios (sales talk, infotainment, conversation, etc.), provides frameworks for various scenarios, and handles the device and media management. Grasbon and Braun describe a generic morphological approach to interactive storytelling in mixed reality environments [5] providing a non-linear interactive storytelling engine [3] based on a rule-based system implemented in JESS [4].

Figure 2 shows different multimodal interfaces integrated into the digital information booth of the IZA project. Exemplarily, microphones are used for speech recognition, a video camera is used for video recognition or a scanner for business cards is used

Figure 2: Digital Information Booth – multimodal interfaces.

to get in contact with the user and to get his email-address in order to provide personalized information.

From the scientific-related point of view, there are several objects of research in Digital Storytelling which lead to new technology for multimedia applications: With new forms of scenario-based design, communication designers shall be included in an interdisciplinary design process for complex applications beyond screen design. Storytelling methods shall be involved to create drafts, design artefacts, and prototypes in a pre-production phase. To include storytelling aspects such as dramaturgy, character, and emotions in an interface, new forms of Conversational Interfaces are proposed. Further, the vision of writing as programming leads to new High-level APIs that allow application programming in a storytelling way.

A current technological trend is the development of intelligent assistance systems. In contrast to a tool that must support direct manipulation and should be handy, an assistant is expected to solve tasks autonomously, in parallel with other tools operated by the user. Proactivity is another criteria: in a museum, a trade show or other kiosk systems, a virtual being as a partner can facilitate orientation in artificial worlds by playing an active role in prompting user interaction. This interaction resembles interpersonal communication much more than interactions using traditional interactive tools.

For this reason, currently anthropomorphous user-interface agents (avatars) within interactive (3D) worlds [11] are developed. Avatars are able to communicate by using human-like interaction techniques, such as lip-sync speech, facial expressions, and gestures. A conversational user interface system should be proactive, narrative, and emotional, and should represent a certain role metaphor. According to the role, not only is the avatar's behaviour and importance important, but also its location and timing.

Application Examples

Apart from digital information booths orinformation kiosks, there are other possible application areas using the same technology of user interaction design, interactive digital storytelling and multimodal user interfaces. For example, within the GEIST project a story engine based on the Propp model [12] is conceptualised as the control unit for dramaturgy [15][6], interactive workflow and further personalized and application-driven parameters in a mobile computer game. From the educational and learning point of view, the main goal of GEIST is to transmit historical information providing mixed reality technology. Users/pupils walk around at historical sites in Heidelberg and get information by virtual avatars presented on head-mounted displays.

In the context of EMBASSI and MAP, users are provided with animated avatars as assistants integrated into entertainment electronics, e.g. presented on a TV screen. Examples for this exist in both multimedia home-entertainment and business applications. In order to improve acceptance of the user interface and to increase immersion, these applications are enhanced by human-like conversation rules for the interaction with user interface agents.

TELEBUDDY represents a physical avatar equipped with different sensors respectively sense organs (see, listen, speak, show) enabling users to contact the Telebuddy via the internet and communicate via different channels for the various senses. For example, the Telebuddy can visit an exhibition and different users (user groups) can look through the eyes of the Telebuddy or speak to other visitors of the exhibition using the Telebuddy speech interface.

With the art-E-fact project (art-E-fact: A Generic Platform for the Creation of Interactive Art Expe-

Figure 3: TELEBUDDY – physical avatar visiting the EXPO exhibition 2000 in Hannover, Germany.

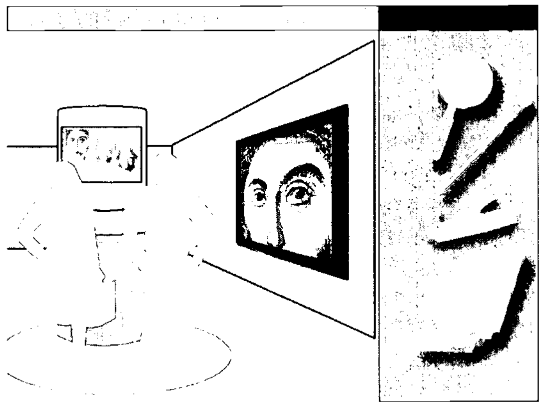

Figure 4: art-E-fact – generic platform

rience in Mixed Reality), an interdisciplinary team of computer scientists, media designers, artists, art historians, restorers and media publishers is reaching for a new stage of convergence of arts and technology. The project will create a platform to develop and explore new forms of creation, presentation and interaction with various kinds of artistic expression. Digital storytelling and Mixed Reality technologies enable new dimensions for artistic expression, and therefore will build the foundation for the new artistic applications.

Therefore, the main objective of the art-E-fact project is to develop a generic platform for interactive storytelling in Mixed Reality to allow artists to create artistic expressions in an original way, within a cultural context between the virtual ("new") and the physical ("traditional") reality. The generic platform for artistic creation in mixed reality is based on a kernel that combines a virtual reality system with a scenario manager for interactive conversations with partially autonomous virtual characters. Further, components for media management and abstract device management enable flexibility and allow authors to design multimodal interactions on various abstraction levels. Results are positive:

- Providing interactive storytelling dialogue structures; instead of "navigation" metaphors in hypertext structures.
- Enabling the design of holistic spectator experiences by integration of design issues concerning content, story, characters, their modalities and used hardware; instead of being constrained to mere screen design with a fixed interaction metaphor.

In summary, it is possible for artists to include anthropomorphic interactions such as speech, gestures, eye movement, and body poses into their design of mixed realities, and to direct lifelike avatars in order to act.

Art-E-fact is a generic platform for the creation of interactive art experiences in mixed reality. Sketched is an installation concept of visitors having conversations with virtual philosophers on the screen (left), with various options for physical props to interact (right).

In the following sections of this paper, the authors describe new methods and concepts to adopt these basic technologies of interactive storytelling and user interaction design and to transmit them into a mobile scenario for museums as one representative of the wide range of edutainment applications. Here, game-oriented concepts are integrated to enhance the usability of developed concepts, respectively user acceptance, especially in the context of pupils and young visitors to museums. The difference between interactive storytelling and gaming [17] is that while both concepts try to attract attendees to walk through a story, within computer games users/players don't have any guarantee of solving the problem, which means it is possible to lose a game. You are always successful within a story.

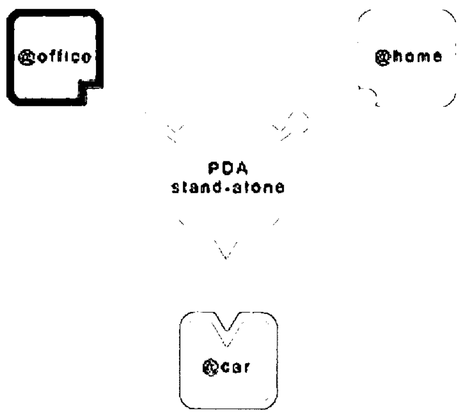

Figure 5:The 4 main areas (incl. periphery) for the Mobile Companion

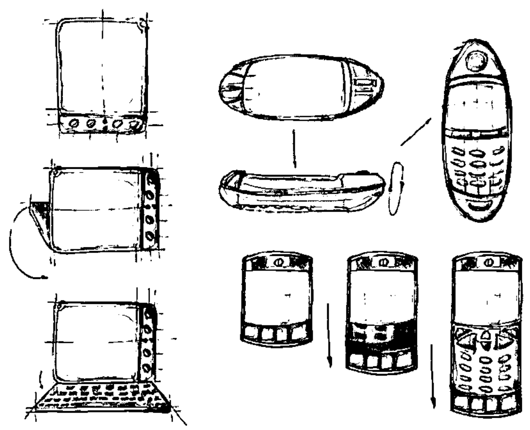

Figure 6: Sketches of example transformations for Smartphones and PDAs

Another difference between story-telling and gaming is that digital story-telling always uses a dramatic concept . Braun,[7], Schneider [1] and Müller et al. [18] describe the usage of conversational user interfaces within these gaming vs. narrative storytelling environments. However currently there is a trend to merge both disciplines and to use storytelling metaphors within narrative environments in the context of cyberspace for computer games [9].

New Methods and Concepts

The Transforming User Interface

The concept of the Transforming User Interface, developed in 2001 by ion2s, tries to solve the problems of the dynamic changes of the mobile situation, context and environment around the user and his own priorities and needs. A digital device, the Mobile Companion, which is in use in four main areas (@home, @office, @car, stand-alone) interacts with the user across the visual, the auditive, and the haptical mode by sounds, images, speech, gestures etc. It transforms its user interface components when it is used in a new area or when the boundary conditions are changing. In the periphery the Mobile Companion and its User Interface components will be prepared for its next tasks.

When it is used in one special field, for example at the office, the interaction principle is optimized for work at a desktop. Priority has the Graphical User Interface. When the Mobile Companion is integrated in the car, becoming a driving assistant system, the locus of the user's attention has to be on the driving process. It is difficult to communicate with the Mobile Companion by the Visual Mode. So the Graphical User Interface must be transformed to an optimized impulse-based interface and the auditive and gesture-based user interfaces become more important. For example, the result of a simple study shows that the changes in the driver's stress influence the Interface transformation. If stress is high, the auditive literfaces or gestures have priority, visual Interaction is very reduced and information concentration must be low. Higher density of content and the interaction per Graphical User Interface is possible in a low stress driving situation. Evaluation and studies of situations, environments, different user interface components and user contexts brought up a need for control equipment to govern the transformations .

These principles and the main concept of the Transforming User Interface can be transferred to every Digital Appliance offering multimodal access and communication methods, and is used in variable and changing environments such as in exhibitions and museums. The present technology enhancements, for example tracking systems, offer the possibilities to understand the context and to focus the human being. This point is very important for building up experience environments for edutainment appliances. Focussing the user and getting him in con-

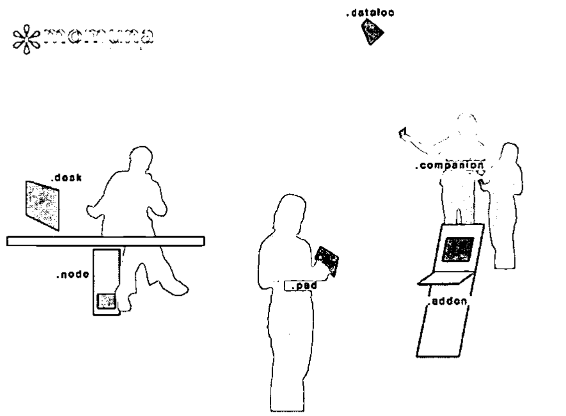

Figure 7: Momuna - scenario at the museum.

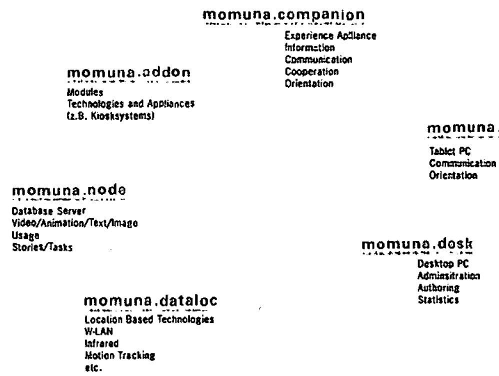

Figure 8: Momuna - components

versation with Applications, Devices and other people enables User Interface Developers to generate innovative digital and interactive edutainment presentations, especially for kids.

Kids Innovation

Kids Innovation is a unit and a method of ion2s to develop user-centred interface concepts for the young target group. In cooperation with specialists for technical questions, storytelling, learning methods, education, pedagogy, etc. new concepts, applications and digital appliances are designed for kids in schools, in museums, in their private areas, etc. From the beginning the young user is completely involved as a full and most important member of the interdisciplinary project team. During the project the kids are organized in workshops for activities and evaluations to deliver useful information for the concepts, about the "needs" and "wants", and regarding demographical and statistical facts of their user group. The transfer of know-how and knowledge lets kids bring up new ways of thinking and is one synergetic effect. Kids Innovation projects help the clients and partners to evaluate their "user kids" during the development process and in the continuous usage of the digital application.

Momuna

A special edutainment concept, to integrate digital experience appliances in a museum, was based on Kids Innovation. The Momuna-concept (mobile museum navigator) is targeted for museums or theme parks, their educational departments and enrichment school programs. It combines simple mobile devices (Momuna.companion/.pad), standard workstations (Momuna.desk), tracking and location technologies (Momuna.dataloc) databases and servers (Momuna.node) and possibilities to integrate any kind of digital presentation devices and technologies (Momuna.addons).

The Momuna-system enables institutions to get valuable feedback about their own exhibition concepts, the behavior of (young) visitors and how they deal with this new digital museum guide. The technologies and additional case studies help to support schools, classes or any other kind of visitor group within the three major phases of visiting a museum: preparation, execution and postprocessing. This sounds very technical but the concept focuses human-computer interaction principles, user interface modes, and users' groups and their needs.

A main digital device, the Momuna.companion, attends every visitor of the exhibition. Communication, cooperation, orientation, user tracking, and location based services are provided by the technological environment installed in the museum. The intelligent use of these technologies, in combination with the mobile devices, enhanced by its game-based and functional concept, expands the museum to an experience learning environment.

The museum, teachers and their pupils are the three main target groups within the Momuna concept. The kids are divided into three more subgroups: the individual, the buddies and the rest. During a story-

Figure 9: IZA mobile – basis for Dino-Hunter using the Momuna concept.

based competition, based on educational aspects, the groups have to cooperate and to communicate via their Momuna.companion among each other. They have to collect relevant information to solve their exercices.All relevant or special highlights can be stored for reinforcement. The Momuna.pad of the teacher offers an overview oft his groups, providing him communication functionalities and other useful features to support his educational work.

Dino-Hunter

The basic methods of interactive digital storytelling, Kids Innovation and the different concepts of user interaction design are combined within the Momuna-based Dino-Hunter, which represents an exemplary demonstrator for an exciting edutainment application in the field of palaeontology exhibitions and museums.The demonstrator shows the possibility to "edutain" pupils providing a story-driven 3D-Puzzle. Experience appliances, complex mixed reality components and storytelling multimodal technologies enable young visitors to become a "Dino-Hunter" during their museum visitation.

Motivated by a basic story, young guests of the palaeontology exhibition get into an experience environment by playing a simulated 3D-Puzzle.The pupils look at artifacts, and identify and collect the virtual fragments of a 3D-Puzzle. In the concept of Dino-Hunter, the parts of the Puzzle are several bones, directly but virtually marked as the dinosaur skeletons. By their digital device and under support of 3D-Tracking technologies they can use their device screen as a window to locate these hotspots. The Momuna companion is used as an archaeological tool within the Dino-Hunter scenario.

In the first phase of Dino-Hunter, all young visitors (e.g. pupils of a school class) have a briefing with group selection and additional information about the simulated time and habitat. All pupils are specialists in finding and collecting bones and fossils.At the central place the group reconstructs a skeleton. Here, the muscles and skin are reconstructed and a complete 3D-model is generated (and optionally animated).Thus, the resulting model is visible from any point of view, using a PDA or any other mobile screening device.

The most important tools to implement the Dino-Hunter scenario are simple personal digital assistants (PDAs).They are used to find, clean, transport and visualize fossil artifacts.The only additional hardware needed is a PC with powerful graphic hardware as central server of the scenario (Momuna.node). The PC works as an information broker and the PDAs are the clients. All clients are continuously tracked (Momuna.dataloc) in order to locate each PDA (group of visitors) at any time.

Figure 10: Triceratops – typical representative of a fossil artifact within paleontological departments of museums.

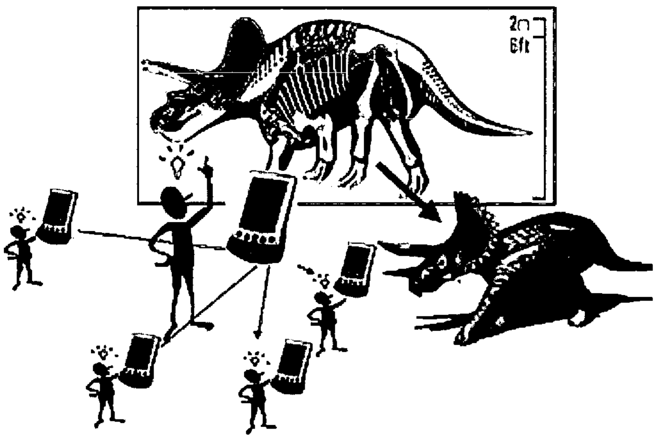

Figure 11: Dino-Hunter scenario – overview and components.

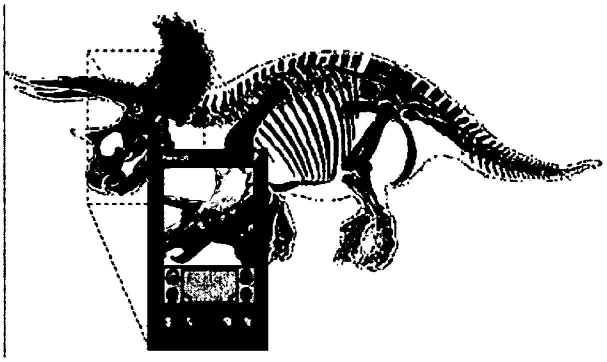

Figure 12: Dino-Hunter – use of augmented reality to show different real and artificial views on virtual triceratops.

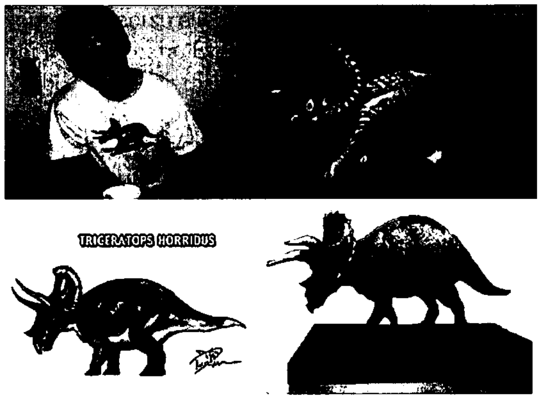

Figure 13: Dino-Hunter – use of augmented reality to show different real and artificial views on virtual triceratops.

Additionally the server tracks and records the viewpoints of each client. Thus each PDA is able to get its virtual camera for real-time visualization later on. To get a more photo realistic view it is possible to use tracked tableau-PCs.

To identify the marked bones the kids move the mobile device of Dino-Hunter across the artifact. The system detects the real justification of the device and tracks the viewpoint of the visitor. Because of the 3D-based tracking technology, the accurate display window, which is behind the digital companion, is able to simulate exactly the scenario/screen. Close to reality the modelled environment of the museum and of the exhibits will be merged with virtual 3D elements. In the concept of Dino-Hunter the mixed reality component can be the marked bones, which the user has to find or a lot of hot spots, which can be additional information but offered by an immersive media presentation. For example, real dinosaur skeletons can be seen "through" the Dino-Hunter´s tool as a real dinosaur full of life. Another example is that you can add organs, muscles, or the skin of the dinosaur and see how it behaves in motion.

Many other scenarios are possible for Momuna and Dino-Hunter. The simplest scenario is called "Walk Around and Understand": Each user has his own PDA, handheld or mobile. He walks around in the museum, and his digital assistant enables him to see all artifacts in the way he wants to. It is possible to see fossils with muscles and inner organs or with coloured skin, fletched or even furred. In another scenario, scientists could compare different styles of reconstruction of extinct animals. Dino-Hunter can reconstruct the way a fossil would walk, run, fly or swim. So even scientists could get advantages from using Dino-Hunter. Further on, Dino-Hunter is easily transferrable to other application scenarios such as mobile games for large booths at an exhibition or any other edutainment scenario such as indoor and outdoor event-tourism.

Summary

This paper presents different interdisciplinary concepts for edutainment applications combining computer science technologies, interactive digital storytelling techniques and Kids Innovation enhanced by user interaction principles. Based on these concepts, Dino-Hunter is introduced as a Momuna-based and game-oriented mobile experience appliance scenario for museums. From the content-related point of view, interactive storytelling techniques are used to establish both an exciting story and narrative environment/information appliance for kids. From the technical point of view, multimodal interfaces and mixed reality platforms are used to create natural, human-like and immersive user interfaces and to improve user-friendliness in general.

The concepts provided could be easily adapted to other application scenarios within the wide range of edutainment applications, e.g. e-Commerce applications such as virtual or physical information kiosks, multimodal exhibition booths, cultural heritage or collaborative learn environments.

Acknowledgements

Methods and concepts provided in this paper have been developed by collaboration between the department of Digital Storytelling at the Center for Computer Graphics and the user interaction specialists of ion2s – buero fuer interaction (engl. ion2s – office of interaction), both located in Darmstadt, Germany. Ion2s is responsible for Kids Innovation, Momuna and Transforming User Interface. The department of digital storytelling is responsible for the technical-oriented and storytelling aspects. Dino-Hunter is introduced as a new concept based on intensive collaboration between the two institutions.

Related Links

ZGDV e.V. Digital Storytelling – *http://www.zgdv.de/distel*

ion2s – buero fuer interaction – *http://www.ion2s.com*

Kids Innovation – *http://www.kidsinnovation.com*

Momuna – *http://Momuna.ion2s.com*

GEIST project, see *http://www.tourgeist.de*

EMBASSI and MAP projects, see *http://www.embassi.de* and *http://www.map21.de*

TELEBUDDY project, see *http://www.telebuddy.de*

art-E-fact project, see *http://www.art-e-fact.org*

References

[1] Braun, N, O.Schneider. Conversation modeling as an abstract user interface component. In: Proceedings of GI Workshop Synergien zwischen Virtueller Realität und Computerspielen: Anforderungen, Design, Technologien, Vienna, September 2001.

[2] Cassell J, S.Prevost, J.Sullivan, E. Churchill. Embodied conversational agents. Cambridge, MA: MIT Press, 2000.

[3] Crawford C. Assumptions underlying the Erasmatron interactive storytelling engine, In: Mateas M, Sengers P, editors. Proceedings of the AAAI Fall Symposium: Narrative Intelligence, Technical Report FS-99-01, Menlo Park, CA: AAAI Press, 1999. p. 112–4.

[4] Friedman-Hill EJ. Jess, The Java Expert System Shell, in SAND98-8206 (revised), Livermore, CA: Sandia National Laboratories, 2001.

[5] Grasbon, D, N.Braun. A Morphological approach to interactive storytelling, In: Fleischmann M, Strauss W, editors. Proceedings: cast01//living in mixed realities. Special issue of netzspannung.org/journal, The magazine for media production and intermedia research, 2001.

[6] Laurel, B.: Computers as Theatre, Reading, MA, 1993.

[7] Mateas, M, A.Stern. Towards integrating plot and character for interactive drama. In: Dautenhahn K, editor, Proceedings of the AAAI Fall Symposium: Socially Intelligent Agents: The Human in the Loop, Technical Report FS-00-04, Menlo Park, CA: AAAI Press, 2000, p. 113–8.

[8] Mateas M, P.Sengers. Narrative intelligence, In: Mateas M, Sengers P, editors. Proceedings of the AAAI Fall Symposium: Narrative Intelligence, Technical Report FS- 99-01, Menlo Park, CA: AAAI Press, 1999. p. 1–10.

[9] Murray J. Hamlet on the holodeck: the future of narrative in cyberspace. Cambridge, MA: MIT Press, 1998.

[10] Penrose, R. (1989).The Emperors New Mind. London (U.K.): Oxford University Pres; p. 398.

[11] Perlin K, A. Goldberg. Improv: a system for scripting interactive actors in virtual worlds. Computer Graphics 1996;24(3):205–16.

[12] Propp V. Morphology of the folktale, International Journal of American Linguistics, Part III. 1958;24(4): 1–106.

[13] Raskin, J. (2000) The Humane Interface – New Directions for Designing Interactive Systems. Massachusetts (USA): Addison-Wesley, ACM Press; p 24ff.

[14] Sengers, P. Narrative intelligence, In: Dautenhahn, K, editor. Human Cognition and Social Agent Technology, Advances in Consciousness Series, Philadelphia, PA: John Benjamins Publishing Company, 2000.

[15] Szilas, N. Interactive drama on computer: beyond linear narrative, In: Mateas, M, P.Sengers, editors. Proceedings of the AAAI Fall Symposium: Narrative Intelligence, Technical Report FS-99-01, Menlo Park, CA: AAAI Press, 1999. p. 150–6.

[16] Spierling, U, D. Grasbon , N.Braun Ilurgel. Setting the scene: playing digital director in interactive storytelling and creation in Computers & Graphics 26, 2002. p.31-44.

[17] Woodcock S. Game AI: The State of the Industry, in Game Developer Magazine, August 2001.

[18] Müller W, U. Spierling, M. Alexa, I. Iurgel. Design Issues for Conversational User Interfaces: Animating and Controlling 3D Faces. In: Proceedings of Avatars 2000, Lausanne, 2000.

Evaluation

Investigating Heuristic Evaluation: A Case Study

Kate Haley Goldman, Institute for Learning Innovation, and Laura Bendoly, Atlanta History Center, USA

http:// www.ilinet.org

Abstract

When museum professionals speak of evaluating a Web site, they primarily mean formative evaluation, and by that they primarily mean testing the usability of the site. In the for-profit world, usability testing is a multi-million dollar industry, while in non-profits we often rely on far too few dollars to do too much. Hence, heuristic evaluation is one of the most popular methods of usability testing in museums. Previous research has shown that the ideal usability evaluation is a mixed-methods approach, using both qualitative and quantitative, expert-focused and user-focused methods. But some within the online museum field have hypothesized that heuristic evaluation alone is sufficient to recognize most usability issues. To date there have been no studies on how reliable or valid heuristic evaluation is for museum Web sites. This is critical if heuristic evaluation is to be used alone rather than in tandem with other methods. This paper will focus on work being done at the Atlanta History Center as a case study for the effectiveness of heuristic evaluation in a museum Web site setting. It is a project currently in the beginning stages of development. The Center is applying a thorough mixed-methods approach to evaluation, including heuristic evaluation. The results of this project will assess how complete and how useful a rigorous heuristic evaluation is alone and in conjunction with other methods in the development and implementation of an online educational resource.

Keywords: Evaluation, Heuristic Evaluation, Usability

The Atlanta History Center as a Case Study

The Atlanta History Center (AHC) has begun a three-year education outreach initiative funded by the Goizueta Foundation to enhance their existing outreach program and Web site and help to develop a first-rate distance learning program. The site in question focuses on the online publication of educational materials and resources developed by the Center for a target population of educators in schools, both classroom teachers and media specialists. Since this population is narrowly defined, yet of prime importance to museums, the project makes an ideal forum for testing heuristic evaluation in a museum setting. The Institute for Learning Innovation has been serving as the evaluator for the Goizueta Foundation distance learning project.

The project has three primary educational objectives, that are to:

1. Improve their existing Web site and develop new content and features so that the site is more easily accessible to educators,
2. Ensure that the content better reflects Georgia's Quality Core Curriculum, and
3. Increase the number of teachers and students who use the Web-based educational materials.

A secondary objective, common to many museums, is to build such strong ties to the educational community that the number of school group visits to the physical site of the Atlanta History Center increases.

While creating the study design for the above-mentioned project, we became concerned about one of the most commonly used techniques in Web site evaluation, known as heuristic evaluation. Heuristic evaluation is a usability engineering methodology where experts who are trained in usability but who are not the end users of the proposed technology project compare the proposed technology against established usability principlesknown as heuristics. The training time for this technique is relatively short- as little as a half-day workshop, and the cost is often lower than other possible usability techniques. Due to this accessibility, heuristic evaluation has been frequently used by museums.

While there are many factors to consider when selecting a research methodology, such as cost, sample size, and personnel, it is assumed that the techniques used must be fundamentally sound. Heuristic evaluation has become hotly debated within the human-computer interaction field due to concerns about the reliability and validity of the results that it produces. Some specialists claim that heuristic evaluation both overlooks usability problems that may cripple the ability of a person to use the program in question, and highlights issues that the user never encounters. Previous research, such as that by Harm and Schweibenz (2001), has shown that the ideal usability evaluation is a mixed-methods approach, using both qualitative and quantitative, expert-focused and user-focused methods. But some within the online museum field have hypothesized that heuristic evaluation alone is sufficient to recognize most usability issues. To date there have been no studies on how reliable or valid heuristic evaluation is for museum Web sites. This is critical if heuristic evaluation is to be used alone rather than in tandem with other methods.

Using the current project at the Atlanta History Center as a case study, we saw an opportunity to further investigate the issue of reliability and validity in using heuristic evaluation for museum Web sites. This paper will outline our proposed techniques and current thinking; as the project develops we expect these techniques to evolve.

Why evaluate at all?

As a point of reference, it is useful to step back from the AHC project and review the goals and methodologies of both traditional museum evaluation and the developing field of museum Web site evaluation. Evaluation is used to urges us to clarify our goals and accomplish our objectives. If we are able to define what we intend to do, we are more likely to achieve our goals, increase the museum's responsiveness to the community, avoid false assumptions about our visitors, and save time and money. Evaluation can be scary, because a project with unclear objectives and no evaluation can always be described as successful. This is perhaps best stated by the Flying Karamazov Brothers who said, "If you don't know where you're going, any road will get you there."

A quick review and comparison of traditional museum evaluation and museum Web site evaluation is covered in Table 1. Audience research is done by some institutions on a regular cyclical basis, by others who have done no other research and need a starting point or by still others who are beginning a new initiative or strategic plan. Audience research provides demographic information and other basic visitor information and is often done on the internet through log files analysis and surveys.

Traditional museum evaluation is made up of four types, not including the above mentioned audience research. Front-end evaluation typically occurs during the initial planning phase of project development and provides information about visitors' interest, expectations, and understanding of proposed topics for a program. Formative evaluation takes place while a project is in development and construction. It provides feedback on the effectiveness of a project and its components — feedback which allows developers to make informed decisions as they continue to build the project. Remedial evaluation is generally conducted after a project is available to the public. This type of evaluation focuses on determining changes which need to be made to the program to improve it. Summative evaluation is conducted after an exhibit or program is completed, and it seeks to determine the extent to which exhibit or program goals were met.

Usability testing is a standard piece of the larger development lifecycle throughout the technology

Museum Evaluation	Museum Web Site Evaluation	Web Site Methodologies
Audience Research	Audience Research	Log Files Analysis, Surveys
Front-End	Usability	User Testing and Think-aloud Protocols, Heuristic Evaluation, Surveys
Formative	Usability	User Testing and Think-aloud Protocols, Heuristic Evaluation, Surveys
Remedial	Usability	User Testing and Think-aloud Protocols, Heuristic Evaluation, Surveys
Summative	??????	??????

Table 1: Evaluation Types and Methods

industry and has been carried over into the field of museum technology development. Usability is currently the main focus for formative, remedial and even front-end evaluation. Although usability is extremely important and is the focus of this current project, the fact that a project or program is usable does not make it de facto valuable, or even used. The logistical and methodological difficulties of assessing the value of a project when the users are geographically scattered means that summative evaluation of museum Web sites is rarely undertaken.

Background on Usability Engineering Techniques

The human-computer interaction field has developed a wide range of techniques to evaluate usability of technology projects. Techniques that are expert-based are known as usability inspection techniques. For-profit companies often choose expert-based methods over user-based methods because of the high costs of doing laboratory tests with end-users.

Heuristic evaluation is one of the most informal methods of usability inspection, meaning it is based on rules of thumb and the skills of the evaluators. In heuristic evaluation, the evaluators may be non-experts who have received some training in usability principles. Since this is a less formal method which avoids using a full set of controls or specified personnel lower costs are incurred than in formal testing. To quote Mack and Neilsen,

> Usability engineering activities are often difficult to justify and carry out in a timely way, but many activities can be done quickly and cheaply, and produce useful results. The methodology decision ...turn less on what is "correct" than on what can be done within development constraints. After all, with sufficient resources we would likely simply aim for rapid prototyping and end-user testing. *(Mack and Nielsen, 1994)*

Although other usability inspection techniques are rarely used in the museum field, we will briefly describe them below in order to give a sense of what could be used or adapted as a technique for our field. The majority of these are designed for designers and developers in the formative development period of a project, rather than in the front-end or remedial stage.

Possible Usability Inspection Methods:

1. Guideline Review

Project is checked to determine conformity to a list of usability guidelines. Comprehensive sets can contain more than a thousand guidelines, and require skilled expertise. They are considered a mix of heuristic evaluation and standards inspections.

2. Standards Inspections

An expert in a particular type of interface inspects the product based on guidelines for that specific product range.

3. Cognitive Walkthroughs

Exploration-focused inspection examines on one feature of usability – the ease of learning. This might be a useful goal for a complex software product, but for a public Web site a more common goal is ease of use. Ease of use would mean a first-time user could navigate and accomplish his or her objective easily, as opposed to finding it easy to become an expert of a more complex system.

4. Pluralistic Walkthroughs

Group meetings with users, developers and human interaction personnel walk through user scenarios, documenting each step of the scenario and discussing implications.

5. Consistency Inspections

Inspections by designers and developers across multiple projects, ensuring that the projects have consistent design elements and usability. For instance, as multiple designers may work on separate functions of a museum Web site, a consistency review would evaluate the congruity of the different sections or how well each section complies with ADA guidelines.

6. Formal Usability Inspections

Inspection method similar to software code inspections. designed to discover and report a large amount of data efficiently. Inspectors take on user roles and work through prescribed scenarios.

7. Feature Inspections

Focuses on whether the project functions as developed meet the needs of the intended end users. In traditional evaluation, this would be a part of summative evaluation.

Reliability and Validity Issues in Heuristic Evaluation

Reliability is the consistency or stability of a measure from one test to the next. Repeated measures of a static item using a reliable measure should end in identical or similar results. Validity is a term used to describe whether a measure accurately measures what it is supposed to measure. For instance, it is hotly debated whether SAT scores accurately assess college achievement. If SATs did accurately assess achievement, they would be a valid measure.

Studies that examine the reliability of inspection methods include two studies by Rolf Molich. In the first study, he asked four commercial usability laboratories to carry out usability tests on a calendar program that was commercially available. One laboratory found as few as 4 problems, another found as many as 98. The biggest concern, however, is that only one problem was found by all four teams and over 90% of the problems found by each team were found by that team alone. The second follow-up study had similar results- there was little inter-rater reliability.

The validity of usability inspection methods should be easier to address – the pertinent question asks: how predictive are these methods of end-user problems? Studies on that question have been completed outside of the museum field. Karat (1994) reports on the results of several such studies. A study by Desurvire (1994) compared heuristic evaluation and an automated cognitive walkthrough to laboratory tests with end users. The system in question was not a Web site, but a telephone system that completed six basic tasks. Table 2 below contrasts the results of the laboratory data with end users and the data collected using inspection methods.

The top line in Table 2 indicates the number of usability problems and interface improvement ideas that were observed during user testing in the laboratory. The remaining part of the table shows the percentage of these problems and improvement ideas found by the evaluators using either heuristic evaluation or cognitive walkthrough. (Desurvire 1994).

In the study above, experts were able to predict at best 44 percent of the usability problems identified by the end users. The table above does not express variance in the problems that occur. Some problems users encounter are relatively minor and others prevent the user from completing major tasks. Desurvire dealt with this issue by asking each participant to assign Problem Severity Codes to the problems uncovered. The table displaying

Method	Evaluators	Problems That Did Occur	Potential Problems	Improvements
Lab	Observed with users	25	29	31
Heuristic Evaluation	Experts	44%	31%	77%
	Software developers	16%	24%	3%
	Nonexperts	8%	3%	6%
Cognitive Walkthrough	Experts	28%	31%	16%
	Software developers	16%	21%	3%
	Nonexperts	8%	7%	6%

Table 2: Prediction Rate of End-User Problems

		Problem Severity Code (PSC)		
Method	**Evaluators**	**Minor Annoyance/ Confusion**	**Problem Caused Error**	**Problem Caused Task Failure**
Lab	Observed with users	5	3	17
Heuristic Evaluation	Experts	80%	67%	29%
	Software developers	40%	0%	12%
	Nonexperts	20%	0%	6%
Cognitive Walkthrough	Experts	40%	67%	18%
	Software developers	0%	0%	12%
	Nonexperts	20%	0%	6%

Table 3: Prediction Rate of End-User Problems by Severity of Problem

these results is reproduced below. Note that experts were able to detect 80% of the minor problems or annoyances but only 29% of the problems that caused task failure.

The Top line in Table 3 indicates the number of usability problems in three severity categories that were observed during user testing in the laboratory. The remaining part of the table shows the percentage of the problems in each of the three categories found by evaluators using either heuristic evaluation or cognitive walkthrough. (Source: Desurvire 1994)

These results raise serious questions about the validity of heuristic evaluation, about the ability of the technique to predict end-user errors. Missing any error that regularly leads to task failure is highly problematic. Worse yet, using heuristic evaluation as the sole usability technique would result in 70% of the errors that cause task failure going undetected in this example. In addition, many interface errors found by the experts using heuristic evaluation are false positives- meaning they find errors that don't actually impact the end-user, wasting development resources on what might not really be a problem.

Still, these results were gathered by a system unlike that used to evaluate museum Web sites. Perhaps the nature of the medium (museum Web sites) allows us to use heuristic evaluation to detect a higher rate of error. Our study aims to replicate this experiment with the AHC Web site.

Research Design for AHC Project

In order to test the reliability of the heuristic evaluation methodology, we will use multiple methodologies, including both heuristic evaluation as well as user testing with think-aloud protocols. These two types of methodology are quite different. Think-alouds are a user-focused methodology where we ask the user to talk-aloud while interacting with the technology, therefore hopefully revealing the conscious cognitive processes of the user. With this technique, the interplay between thought and action is revealed by the user, rather than assumed by the researcher.

Within usability engineering, an iterative design structure is critical, and the most complete designs incorporate a cyclical process of inspection methods and user testing at different points within the evaluation process. This allows a set of checks so that the solution to an interface problem does not create increased errors in other functions. For the purposes of this experiment, each technique will be performed on the exact same version of the Web site. (In a typical design structure, end-user testing would occur after changes from the heuristic evaluation had already been incorporated into the Web site.) For AHC project itself, there will be several iterations of evaluation that are not a part of this experiment.

In each of the methodologies used, we will develop scenarios or tasks for the experts or end-users to complete. There are advantages and disadvantages to using the scenario approach. If carefully constructed, the scenarios can assist participants in focusing their efforts on specific interface elements. On the other hand, facilitating a more open-ended inquiry will emulate the way most users experience a site, through intuitive exploration. Testers will usually then form their own scenarios with which to make sense of a site. Given that the AHC project is only one piece of a much larger site, we opt to control the scenarios. Complexity of the scenario can at times change the usability issues found, but as the interface here will be fairly straightforwardly task oriented, we do not anticipate this to be a mitigating factor.

Below we will lay out the specific processes for each methodology.

Heuristic Evaluation

The first step in heuristic evaluation is to decide which set of heuristic principles to use. There are many different types of usability principles. Some of the standard ones were developed by Neilsen and others in the early 1990s. (See Tables 4 & 5) By combining the principles from several different sets, we will develop a set of usability heuristics for the AHC project.

For the actual process we will recruit 6 evaluators. Some studies show a benefit to evaluators working in teams, while other studies show a concern that teams "filter out" valid issues. To reap the most benefit, two evaluators will work together while the rest will work individually. Evaluators will be museum professionals who are unrelated to the project at the Atlanta History Center. In order to test the "quick-to-learn" claim of heuristic evaluation, we will not be usability experts. (There is no certification for the usability profession at this time. Within the field, the expert status normally is seen as obtained after 7 years in the field.)

Since the evaluators will not be usability experts, but museum professionals, training will first be given on heuristic evaluation, including both the process and the specific principles for this evaluation. Evaluators will not be familiar with the system itself and may or may not be familiar with the proposed types of users (generally classroom teachers, but also possibly media coordinators and students), types of tasks that system users will be trying, and the contexts involved. Training will be provided to try to put the evaluator into the users' shoes. Evaluators will then be asked to imagine several scenarios while using the site. All scenarios will be described without screen-shots or specificities that would bias the evaluator in how they might approach the site. Evaluators will have an hour or more to complete the evaluation, and will be asked to resist discussing their results with others while moving through the scenarios. We will suggest that evaluators complete each scenario twice, once to gather a rough idea of the problems, and then a second time to link those problems specifically to the defined heuristic principles. Evaluators will be asked to describe in writing each of the specific issues that arise.

After the formal evaluation, a debriefing session will be held to discuss the characteristics of the site, and identify any possible alternate approaches if critical issues arise. After the brainstorming session, evaluators will be asked to rate the severity of the problems they encounter. Severity rating assists developers to prioritize the changes needed in a project.

Simple and natural dialogue
Speak the users' language
Minimize the users' memory load
Consistency
Feedback
Clearly marked exits
Shortcuts
Precise and constructive error messages
Prevent errors
Help and documentation

Table 4: Example of Usability Principles by Molich and Neilsen (1990)

Visibility of system status
Match between system and the real world
Use control and freedom
Consistency and standards
Error prevention
Recognition rather than recall
Flexibility and efficiency of use
Aesthetic and minimalist design
Help users recognize, diagnose and recover from errors
Help and documentation

Table 5: Example of Usability principles by Neilsen (1994)

Neilsen's severity rating is made up of three factors:

1 = The frequency with which the problem occurs: Is it common or rare?

2 = The impact of the problem if it occurs: Will it be easy or difficult for the users to overcome?

3 = The persistence of the problem: Is it a one-time problem that users can overcome once they know about it, or will users repeatedly be bothered by the problem?

Neilsen also mentions a fourth factor which he does not directly add to the others – one of market impact. He points out that certain types of usability problems can have a "devastating effect" on the a project, even if the problem is supposedly easy to overcome.

We will use an alternative system by Desurvire (1994) for severity ratings, one which splits the ratings phase into two different three point scales. The first scale, the Problem Severity Code (PSC), rates the error severity as follows:

1 = Minor annoyance or confusion

2 = Problem caused error

3 = Problem caused task failure

The second scale measures the attitude of the user towards the system, an extremely important variable in the likelihood of a user's continuing with a system once errors have occurred. The ratings for this scale are below:

1 = Content with the system

2 = Frustrated with the system

3 = Exasperated with the system

At times it is difficult to get useful severity estimates from evaluators during the actual session, when they are mainly focused on the finding of problems, rather than on the severity of the problem and how that particular problem impedes the overall purpose of the project. His suggestion is to ask the evaluators to revisit their list of problems after the debriefing session, despite the fact that the evaluators would generally not have access to the system in question.

After gathering the severity ratings, we will do several tests of inter-rater reliability, including calculating the average correlation between the severity rating provided by any two evaluators, using Kendall's coefficient of concordance, and we will also estimate the reliability of the combined judgements by using the Spearman-Brown formula.

End-User Testing

To contrast with the heuristic evaluation, we will complete a round of user testing at the same point in the formative development process of the Web site. We will attempt to have a minimum of 15-20 user-testing sessions. Unlike in the heuristic evaluation phase, users will work separately under the assumption that most end-users of the AHC site will be working on their own. Sessions will take place either in the History Center classrooms or within a usability laboratory. Users will be recruited through the large teacher network that has worked previously with the Atlanta History Center.

Users will be given a series of tasks and asked to work through each of them while articulating their thoughts out loud in a stream-of consciousness fashion. As with the heuristic evaluation phase, users will interact directly with the interface. With each user will be an observer/facilitator who will record users' thoughts and actions as well as use appropriate prompts to probe for further information. Sessions will be audio taped and /or videotaped for further analysis.

Analysis

During both phases of testing, data will be collected on variables task completion, error data, time to complete task, error severity, and user's attitude (the PSC and PAS scales mentioned above) based on the observation of and discussion with the end user. We will provide analysis similar to Desurvire's, doing a comparison of heuristic evaluation and end-user testing on each variable. We will also present analysis on which heuristics are cited most often. If possible, we will present a comparison on the use of evaluators individually and in teams. Finally, we will present recommendations for the use of heuristic evaluation to inspect museum Web sites as well as suggestions for future research in this field.

References

Bailey, B. (2001) How reliable is usability performance testing? Last updated Sept 2001. Consulted August 27,2001. http://www.humanfactors.com/downloads/sep012.htm

Desurvire, H. (1994). Faster, Cheaper!! Are Usability Inspection Methods as Effective as Empirical Testing?. In J. Nielsen and R. Mack (Ed.) *Usability Inspection Methods*.. New York: Wiley & Sons, Inc, 173-199.

Di Blas, N., Pai Guermand, M., & P. Paolini (2002) Evaluating the Features of Museum Websites. In D. Bearman & J. Trant (Eds.) *Museums and the Web 2002 Proceedings*. CD ROM. Archives & Museum Informatics, 2002.

Harm, I. & W. Schweibenz (2001) Evaluating the Usability of a Museum Web Site. In D. Bearman & J. Trant (Eds.) *Museums and the Web 2001 Proceedings*. CD ROM. Archives & Museum Informatics, 2001.

Karat, C., (1994). A Comparison of User Interface Evaluation Methods. In J. Nielsen and R. Mack (Ed.) *Usability Inspection Methods*.. New York: Wiley & Sons, Inc, 203-230.

Mack, R. & J. Nielsen, (1994). Executive Summary. In J. Nielsen and R. Mack (Ed.) *Usability Inspection Methods*.. New York: Wiley & Sons, Inc, 1-23.

Nielsen, J., (1994). Heuristic Evaluation. In J. Nielsen and R. Mack (Ed.) *Usability Inspection Methods*.. New York: Wiley & Sons, Inc, 25-61.

Wharton, C., Rieman, J. Lewis, C. & P. Polson, (1994). The Cognitive Walkthrough Method: A Practioner's Guide. In J. Nielsen and R. Mack (Ed.) *Usability Inspection Methods*.. New York: Wiley & Sons, Inc, 105-139.

New Vision, New Realities: Methodology and Mission in Developing Interactive Videoconferencing Programming

Dr. Patricia Barbanell, Dr. John Falco, Schenectady City School District, and Dr. Diana Newman, State University at Albany, USA

http://www.projectview.org

Abstract

As museums throughout the world enter the interactive arena of digital communications, a need has emerged to access strategies of program development that seamlessly interface with existing missions and resources. This paper describes how Project VIEW, a US Department of Education Technology Innovation Challenge Grant, collaborates with major museums - the Guggenheim (NY) Museum, the Whitney Museum, the Philadelphia Museum of Art, and others - to create templates for developing point-to-point interactive videoconferences with asynchronous Web-based resources that enhance student learning. Underlying the work of VIEW is the premise that, to achieve the promise of interactive technologies, it is necessary to change educational pedagogy. To accomplish this, Project VIEW employs multi-phase integration techniques that bring together the needs and missions of diverse yet intersecting educational delivery systems at museums and schools. While the development of a model for sustainable program development and content integration was a core component of Project VIEW, the primary goal of VIEW is to deliver instruction that produces evidence of higher student learning and academic performance. Interim evaluation by an external reviewer (The Evaluation Consortium, University at Albany, State University of New York) indicates evidence of enhanced student learning among students who participate in Project VIEW programs. Importantly, evidence confirms that outcomes for students are the result of VIEW training and development processes in which both schools and museums are transforming the way that they deliver education and through integrated, interactive videoconferencing and Web-based learning.

Keywords: Interactive Videoconferencing; Pedagogy; Collaborative Development; Constructivist Activities; Mission alignment

Introduction

Project VIEW, based in the Schenectady City School District, NY, is midway through a federal technology innovative challenge grant that is changing the way that Museum Educators and Schools interface in the classroom. To accomplish this, Project VIEW has effectively created a consortium of over 500 demographically varied school districts and diverse, content-rich education providers from across the nation who are pursuing a shared common goal: to develop and use interactive educational programs that expand traditional school communities and bring unique resources to students who may not otherwise have access to them.

From its inception, Project VIEW has recognized the critical need to prepare students with the high degree of literacy needed to succeed in the 21st Century information age (Project View, 1999). To meet this need, Project VIEW has built a constructivist inquiry-based environment for developing innovative educational resources.

VIEW utilizes the unique expertise of both museums and participating teachers to create content which facilitates seamless integration of video programming and supporting Web-based resources into curriculum delivery. Participating museum educators and teachers are merging their educational methodologies into a student-centered, constructivist reality in which literacy is built on multi-media foundations, and individual learners create meaning from digital content formats. (Meuwissen & Geehan, 2001) What follows in this paper is a brief overview of the VIEW model and preliminary data on the effectiveness of the model for both schools and museums.

Phase I - Establishing A Foundation For Program Development.

There are multiple components that are essential for engaging in the VIEW Development process.

These components are not linear but rather occur in a constructivist environment, emerging both in real time and asynchronously, as needed.

Program Prerequisites

VIEW has identified four essential components needed for the development of all programs – collaborating teams, multi-disciplinary instructional expertise, internal institutional commitment, and content standards alignment. These components are constructed on the firm foundation of evidence and experience documented in the Schenectady City School District formal program evaluations.(Evaluation Committee, Title III, 1997-2002)

- ***Collaborating teams:*** In developing programs, study has shown that collaborating building-based teams of at least 3 or 4 teachers are required to create a sustainable community of users. Three or four teachers in a single location can support each other and disseminate the program to their peers. Further, they are likely to continue to engage with a program with energy even if one team member moves on to another location.

- ***Multi-disciplinary instructional expertise:*** Including multiple curriculum areas in building a team is an important requisite for program success because our experience has shown that such content diversity creates a breadth of vision and usability in programs that facilitates sustainability. In addition, the curriculum products and models created by multi-disciplinary teams have broader and deeper applicability for use in a variety of settings.

- ***Content provider institutional commitment and synergy:*** It its essential that development of interactive videoconferencing programs and supporting content be aligned with the mission and expertise of the content provider. Without this important requirement, the program will have no roots or continuity within the institution, and would most likely lack the content underpinning which forms the core of a successful project.

- ***Content standards alignment:*** Alignment with content standards is essential if a program is to offer substance that supports and expands classroom mission and practice. Furthermore, in order for schools to continue to participate in programs, the programs must meet mandated educational needs.

Pre-Development Preparation and Planning:

Before proceeding with development, VIEW requires a series of collaborative planning sessions to ensure that an adequate foundation for development exists. These sessions are designed to create a basis through which a program can be constructed. Focus for such sessions (as needed) may include:

- ***Capacity and connectivity of the provider and/ or participating classrooms:*** The core of the VIEW model is built on a philosophical and pedagogical foundation for interactive and sustainable interactive videoconferencing and Web usage. In designing programs that facilitate engagement with interactive technology, it is essential for all partners to have a baseline technological capacity (hardware, software and connectivity) to fully participate in the development process. To ensure that the infrastructure is in place for all participants, Project VIEW staff meets separately, through a Collaborative Visioning process, with provider and school technical staff to identify equipment and connectivity needs, and to fill baseline gaps where they are identified.

- ***Training and experience of participating provider and school staff:*** While VIEW programs are created so that they can be expanded and/or solidified as technology and technological skills improve, it is crucial that all participants in development have an essential baseline of training and experience. Such basic technological literacy ensures that participants can fully contribute to the interactive process of creating content and delivery modes. To enable hands-on participatory development, VIEW facilitates a constructivist Collaborative Visioning process with participating teachers and museum staff. Technological learning tutorials are offered, as needed, to provide appropriate skills for content development. This process creates a self-motivated assessment and understanding of capacity and readiness among participants.

- ***Content focus aligned with provider institution mission, goals and collections:*** Prior to beginning the collaborative development, Project VIEW staff meets with the museum to establish a program content focus that is aligned with the provider institution. Programmatic options and preferences of the institution are reviewed, and a general focus for the program is collaboratively chosen. Target grade levels and possible curriculum specialties appropriate for the content are also suggested.

It is essential that the content of the program be aligned with collections. Without this foundation, the program would be an empty shell, failing to utilize the unique capacities of the museum. In addition, it is necessary that content be aligned with institution mission and goals to make certain that it can be sustained over time. Such alignment helps to make it likely that the expertise and focus of the institutional programming will continue in the program content area and facilitate sustainability.

- ***Selection of teacher team members aligned with content focus:*** Project VIEW collaborates with schools to select participating teacher teams based on content focus, target grade levels and possible curriculum specialties. In this process, Project VIEW records suggestions from schools regarding content appropriateness, standards interface and teaching applicability. If necessary, Project VIEW meets with content provider to collaboratively consider comments and make adjustments to initial content focus, as appropriate.

- ***Pre-program orientation of teachers and/or provider staff:*** Pre-development orientation and/or introduction sessions are offered separately to participating teacher teams and provider staff. As needed, Project VIEW brings introductory presentations and videoconference demonstrations to participating sites to help establish the baseline of programmatic understanding needed to adequately participate in the development process.

Phase 2 – Exploring Content And Developing A Program

Developing project content and ensuring that the content is useful in enhancing education delivery are among the central activities of Project VIEW. There are multiple components to these activities, and most of them occur in a sequential timeframe.

Content Immersion

Teachers and museum educators joining VIEW collaborations are expected to use their unique and diverse expertise in VIEW program development. Specifically, teachers are expected to be leaders in the area of classroom curriculum and school-based academic pedagogy, while museum personnel are expected to be specialists in the specific content areas of museum collections and archives (Note, (while this paper refers primarily to Museums and Museum education, Project VIEW has done its work with a spectrum of content providers, including Museums, Zoos, Historical Societies and Libraries). Yet, as the Project VIEW model has developed, it has become apparent that when there is a constructivist blending and/or interface of the competencies of the participants and participating communities, the development of sustainable programming is more effective.

The constructivist blending is based on a theory that meaning is more profound when it is constructed, not transmitted. Participants use their own expertise to collaborate with each other to construct methods and models based on intersecting missions and competencies.

To facilitate this process to the greatest degree for each project, the VIEW program development begins at the provider institution with a full day session of content immersion. Participating teachers receive an introduction to collection artifacts and resources from the content provider staff as well as an overview of the proposed program content focus. General responses and discussion of the applicability and utility of the content for classroom use take place, and a basic framework for the program development is discussed. Teachers also reflect on grade-appropriateness of the proposed content focus, and relevance to learning standards and the curriculum.

Development Sessions

Constructivist Activities:

Within two weeks of the ***Content Immersion*** session, participants come together again at the Project

VIEW Laboratory to engage in a constructivist process that contributes to development of a replicable videoconference program with supporting resources and activities. Collaboratively, participants review the ***Content Immersion*** experience and begin to select key concepts that can expand and enhance student learning and achievement. Often during this phase of the development, the content focus narrows, and tangential areas of exploration are tagged for program enrichment and/or exploratory study. Throughout this process, it is important to remember that "Different perspective, interpretations and criticism must be shared and creative conflicts (that lead to new discourse and new knowledge) must be engendered." (Jeffers 2003)

The process of creating the project varies from program to program, but a template for effective development methodology is emerging.

In its developing form the VIEW model includes the following:

- Facilitated constructivist discussion of all participant points of view on content applicability and presentation;
- Structural outline of videoconference presentation and identification of experiences and resources needed to support and deliver the content;
- Self evaluation of program and process and self assignment of specific tasks for independent work.

At the conclusion of the first three development sessions, participants identify areas of content that are needed for program completion, and they collaboratively designate team members to gather and/or create that content.

Independent development by teachers and museums

Following the content immersion and development sessions, participants and museum specialists spend the next 8 to 10 weeks pursuing their team designated development tasks. These tasks vary with the competencies and interest of the participants. In general, museum staff focuses on the collaboratively determined program content, gathering and sharing resources with participants, and authoring a videoconference presentation that enhances understanding of some key aspect(s) of the content. Teachers construct pre- and post-visit activities and resources to support the anticipated videoconference. All participants are encouraged to use technology to interact digitally with each other and to share and pilot parts of the program as they are developed.

Reflection and Completion of Development

After the independent component of the collaboration, the teams are reassembled for a final two days. They conduct a review of their collective efforts including revisiting their goals and evaluating the results of their independent work (curriculum, resources, lesson plans, etc.). The content provider presents the videoconference that they have created, and teachers critique it. Gaps in the program are identified, and the group collectively works to fill them as appropriate. VIEW criteria for high quality programs are emerging as a result of this process. Benchmarks for program excellence are currently being created. However, early review suggests the following may central be among them:

- Seamless curriculum interface;
- Effective use of interactive media;
- Strong use of supporting interactive technologies;
- Alignment of program elements with institutions missions.

At the conclusion of the two days, participants schedule test dates with the provider for field testing the program in the classroom.

Phase 3 – Field Testing of Program

An important component of Project VIEW is to provide an educational reality for laboratory field testing of programs that are developed. This occurs in two stages.

Field Testing with Participating Program Developers

Following the development phase, programs are piloted in the schools of the teachers who helped develop them. Imbedded in those pilots is a commitment among all participants to provide constructive feedback on the effectiveness of the program in the classroom with students. After responses are gathered, the provider is charged with retooling and/or revising programs as needed.

Field Testing with VIEW Trainees

The second stage of field testing takes place in the context of training sessions conducted by project VIEW staff to prepare teachers to be users and consumers of the programs that are developed. As teachers learn how to design integrated classroom modules using interactive video programs and technologies, they are charged with piloting the programs developed by VIEW collaborations. Their commitment includes developing integration plans for their classrooms and providing candid responses regarding the effectiveness and utility of the programs in the classroom. The feedback is shared with the providers and the process of improving programs continues.

Phase 4 - Assessing Progress And Moving Forward

Key to reaching and validating Project VIEW goals is the constant reflection on and realignment of components of the project. To engage in this process, Project View is involved in a continuing multiphase, mixed-methodology evaluation which includes both formative and summative evaluation constructs.

Formative Evaluation

Part of the evaluation is integrally woven into the evolution of the VIEW models on an on-going basis. Reflections and findings of the evaluators are discussed and integrated into evolution of the development models and program structures. Reflections on methodology and content integration occur on a regular basis and have resulted in the following outcomes (Newman, D.L., Smith, J. & Geehan, M.M., 2000).

- More effective organization of development model procedures to allow for different teaching styles of participants including structural integration of interactive collaborative planning and delivery of programs and mission alignment between school and museum educators.

- More effective documentation of curriculum projects through a better organized Web representation that includes pre- and post- program activities and resources to enhance videoconference content.

Summative Evaluation

Summative evaluation has helped to create benchmarks of achievement and signposts of productive directions. Among the trends beginning to emerge are the following findings:

- Increasingly, museum personnel appreciate videoconferencing as not something that is in competition with the museum, but something that supports current museum efforts.

- Participating schools perceive that educational videoconferencing serves to benefit students in their classrooms by increasing their access to authentic and exceptional resources, from historic and scientifically unique locales and classrooms in other areas to experts in various fields and celebrated artifacts.

Concluding Comments

VIEW's emerging model of collaborative development is changing the way that both museums and schools deliver content in the classroom. By creating a process to merge discrete missions and goals of participating institutions, Project VIEW has achieved success in designing and modeling key components required to establish a productive interactive multimedia collaboration among museums and schools. The constraints of schoolhouse walls are melting away as programs reach out to bring museum expertise into the classroom through interactive videoconferencing.

Using interactive videoconference programs and supporting Web-based resources, VIEW makes new and innovative use of museum expertise and collections while structuring content delivery so that it can be seamlessly integrated into curriculum delivery. Museums are realigning their content presentation to interface with and enhance classroom curriculum presentation and content requirements. The result is a new, emerging reality in classroom pedagogy and a new, exciting extension of interpretive activities of museums.

References

The Evaluation Consortium, Evaluation reports of New York State Title III Literacy Challenge Grants, Schenectady City School District, NY (1997-2002).

Carole Jeffers, *Art Education*, Jan. 2003, p. 21

Innovation Challenge Grant 1999

Meuwissen, K.W. & Geehan, M.M. (2001). Creating a Model for Educational Videoconferencing: Three Approaches to Developing Videoconference Products and Related Curricula. A paper at the 15th Annual Edward F. Kelly Evaluation Conference, Toronto, Canada, 2001.

Newman, D.L., Smith, J. & Geehan, M.M.,(2000), Systems change in the human service domain: A model for documenting and evaluating change. A paper presented at the 2000 American Evaluation Association Annual Meeting Honolulu, HI.

New York State Title III Literacy Challenge Grants, Schenectady City School District, NY (1997-2002).

Project View, Grant Proposal, "Proposal Significance" p.3-4 Project No. R303A000002, Federal Technology.

A Rolling Evaluation Gathers No Moss

Lee Anne Burrough, DuPage Children's Museum, Lorrie Beaumont, Independent Evaluation Consultant, David Schaller and Ethalinda Cannon, Educational Web Adventures, U.S.A.

Abstract

Continuous, formative evaluation, from inception to completion, with rapid revisions, was key to the successful development of Kids Design Network (KDN) by the DuPage Children's Museum (DCM).Although a thorough evaluation plan was developed at the beginning of the project, it soon became apparent that additional strategies were necessary to insure success. As new issues were identified and planned techniques were found inadequate, new evaluation techniques were developed and implemented.The development of KDN illustrates the importance of flexible, formative evaluation from the inception of a program to its completion. KDN is an interactive Internet-based learning program for elementary aged students. Using the KDN Web site (http://www.dupagechildrensmuseum.org/kdn), students work collaboratively with an engineer (via real-time communication on the KDN Engineer Chatboard), their classmates and teacher to design, build and test a gadget that meets an engineering challenge. Interactive portions of the Web site include a drawing program where students draw a design for their gadget and a whiteboard with chat where children and engineers can mark on the student's design and communicate via text chat. DCM brought in an evaluation consultant in the very early stages of KDN development.The evaluator became an integral part of development and continuously collaborated with the team.The collaborative nature of the evaluation and the fluidity of the plan ultimately produced a workable, user friendly Web site consistent with the museum's mission.

Keywords: evaluation, interactive Web site, elementary schools, consult with professional, hands-on activities, engineering

Introduction

The Internet's potential for linking students to experts in a field has long been recognized. Many museums and government agencies (including the DuPage Children's Museum) have outreach programs based on a Web site where information is presented by the institution and/or the students, and students and experts communicate via e-mail. Kids Design Network (KDN) takes the important next step and makes this two-way communication *live*, *secure*, and *integral to the learning process*. The DuPage Children's Museum (DCM) conceived KDN as a way to encourage and facilitate the incorporation of design technology, open-ended problem solving, and engineering and computer skills into elementary school classrooms. Basing KDN on the Web extended the Museum's reach well beyond the Chicago area.

Working with young children as DCM does has shown the importance of continuous evaluation of new programs in scenarios as closely approximating actual use as possible. When one combines the unpredictability of children with the wide variety and constantly changing nature of classroom settings and computer technology, flexibility in this continuous evaluation becomes of paramount importance. Continuous, formative evaluation with rapid revisions from inception to completion was key to the successful development of KDN.

The basis for KDN is DCM's Kids Design Engineering (KDE) Learning Lab in which students are challenged to design, build and test an air (balloon) powered vehicle. Now in its seventh year, KDE has been thoroughly tested and is highly successful at teaching elementary students design engineering and problem solving skills. KDE is heavily facilitated by DCM personnel who visit more than 60 classrooms each year, including those in the local school districts where KDE is incorporated into the 3rd grade curriculum.The plan for KDN was to take the KDE model and substitute an on-line engineer for the DCM facilitators. Using the Internet, students could communicate in real-time with the engineer, asking him/her questions about the design of their projects. Based on this idea, DCM applied for and received KDN pilot funding from the Illinois State Board of Education Center for Scientific Literacy.

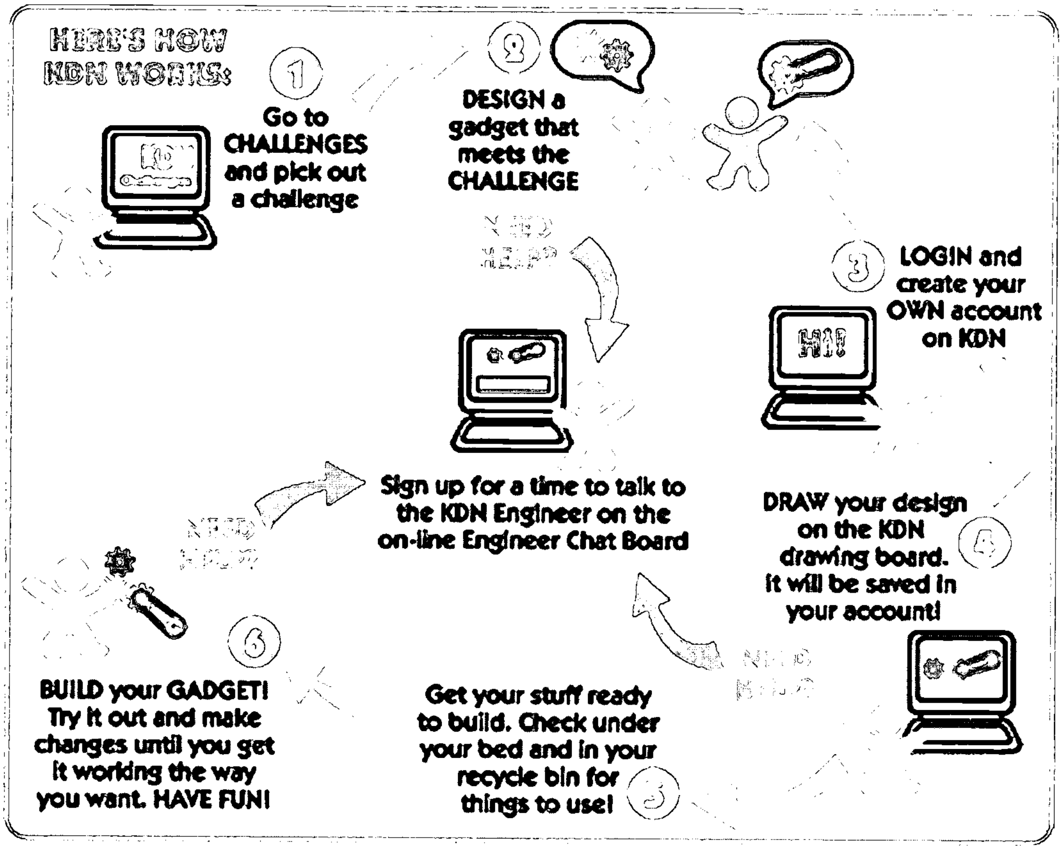

Figure 1: How KDN works.

Program Description

Students can consult with the on-line engineer at any time during the design and building process. Many teachers have their students talk to the engineer more than once. The real-time communication provides students with usable information as they design and build their gadgets.

On the KDN Web site (http://www.dupagechildrensmuseum.org/kdn), engineers introduce a challenge and relate it to a real-life engineering challenge they have faced. Table 1 summarizes current challenges. Elementary aged students work collaboratively with an engineer (via real-time communication on the KDN Engineer Chatboard), their classmates and teacher to design a gadget that meets the challenge. Students use KDN's on-line drawing program to draw a design for their gadget. On the KDN Engineer Chatboard, both students and engineer see the design, mark on it, and discuss it via live text chat. Students build their gadget using materials found in their recycle bins, then test and modify it with help from the on-line engineer, teacher and classmates.

Evaluation philosophy of DCM

All programs and exhibits are continuously evaluated at DCM. Revisions are made based on the evaluation, and then are themselves evaluated. No matter how small the project or how tight the budget, evaluation starts at the first brainstorming session and never really ends. In developing both KDN itself and its evaluation plan, we found that we were following the engineering method that KDN is designed to teach: design, build, test, redesign, build, test again, redesign… . In a project as complex as KDN turned out to be, ongoing evaluation of both macro and micro components was imperative.

Challenge	Task
Puppet	Build a puppet that moves in at least 2 ways
Chemical engineering	Build a gadget to sort 3 different types of objects, such as paper clips, marbles and crayons
Slowly Moving Marble	Build a pathway for a marble to roll down as slowly as possible
Nuclear Engineering	Build a gadget to retrieve something from about 3 feet away from you
Door Alarm	Build an alarm for your bedroom door
Civil Engineering	Build a gadget that allows you to pour water from a jug without touching the jug

Table 1. Engineering challenges currently posted on the KDN Web site.

This philosophy continues to permeate the development of KDN, and typically includes professionals from many different backgrounds. Groups involved in evaluating KDN included:

- In-house members of the KDN team and other DCM employees;
- An Advisory Board drawn from the community and consisting of engineers; computer experts; experts in Web-based school-university-partnerships; and school district administrators, principals, teachers and curriculum coordinators;
- Students and teachers using KDN;
- Web designers and programmers contracted to produce the Web site;
- Other professional Web designers;
- An independent, professional evaluator.

The KDN team consisted of the DCM project manager and experienced KDE facilitators (including engineers, scientists, and former classroom teachers), the independent evaluator, and the contracted Web site developer's programmer and project director. This group worked closely together for 2 years to develop KDN.

DCM Employees

In-house members of the KDN team included former classroom teachers, engineers, scientists, and exhibit and program designers. They were experienced in conducting effective science programming for elementary children, as well as expert in the development of open-ended problem-solving activities. Other DCM employees involved included facilitators of school and museum floor programs, art specialists, exhibit developers, and DCM management. They provided fresh eyes to review all aspects of KDN, and were a ready source of Web site testers. Since these people were not involved in the day-to-day development, they provided a different perspective on everything from Web site graphics to navigation to teacher guidance to ensure KDN was consistent with DCM's mission.

Advisory Board

DCM regularly uses community-based advisory boards when developing new programs and exhibits. KDN's advisory board was especially important in providing philosophical and practical guidance in the early stages of development. School officials on the advisory board helped insure that KDN would be useful to and usable by classroom teachers. Principals and curriculum coordinators helped pilot test KDN in their school's classrooms. Computer experts gave advice on available technology. Experts in school-university-professional partnerships provided information on existing collaborative networks, how these collaborations worked, and potential pitfalls KDN might face. Engineers helped provide Challenges and advice on how to staff the on-line engineer positions. Since engineers would be using the Web site, their input on the components necessary for effective communication were invaluable.

Students and Teachers

It is part of the culture of DCM that all programs and exhibits are pilot tested first in small groups and then in the setting where they will be used. Feedback from the users is absolutely necessary to successful development. Because teachers are the primary users of KDN, a strong emphasis was placed on meeting their needs. Our evaluation methods needed to identify teacher concerns and ways to alleviate them. We found that teachers' opinions/reactions to KDN after using it with children were markedly different from their reactions before using it. For example, teachers greatly underestimated students' enthusiasm for talking with the engineer and the value of the building process. We therefore tested KDN in classrooms and computer labs with entire classes of children. We held preliminary briefing sessions with teachers, observed and talked with them while their students were using KDN, and whenever possible, talked with them afterwards. Ultimately, 40 classroom teachers participated in the development of KDN. Their input was a critical factor in development.

We knew in advance that elementary teachers have very little discretionary time and are under constant pressure to demonstrate that their students perform well on standardized tests. We needed to ensure that teachers could see the benefits of KDN in achieving educational goals. We also needed to ensure teachers were given enough information to feel comfortable using KDN on their own. Many teachers are not accustomed to using hands-on activities in their classrooms, and many are not adept at using computers.

Pilot testing with children is an absolute necessity because of their widely varying skill levels, interests, and experiences. Sometimes our hypotheses were confirmed in the tests, but frequently we were surprised. Testing with large numbers of children is critical for a program intended for use in classrooms. Children function differently in large classroom settings than they do in small groups in or out of school. Tests in schools also gave important information on Web site performance on school computers with many children (up to 30) using the site's interactive features at once. Every part of KDN was extensively child tested. with over 900 children eventually involved.

Web site Developers and Programmers

When KDN development began, DCM had no in-house Web site development expertise. Many potential contractors were considered, with emphasis on identifying one which could provide interactive, innovative and engaging programming for children. Educational Web Adventures, LLP (Eduweb) was chosen. Eduweb worked closely with other members of the KDN team in developing the evaluation plan as well as the Web site. In addition to normal extensive preliminary testing of Web site components, Eduweb's programmer and project director were in the classroom observing children for the first full-scale test of the KDN. Subsequently, they frequently served as on-line engineers, getting first hand understanding of how the Engineer chatboard works. The programmer continues to periodically monitor engineer chat sessions and trouble-shoot as needed.

Independent Evaluator

When given the choice, many museums conduct front-end and formative evaluation in-house and save their budgets to bring in an independent evaluator at the summative stage when the project is complete. DCM brought in an independent evaluator at the very early stages of KDN development. This designated one person as the focus for evaluation, and the evaluator became an integral part of the KDN team. She helped define the evaluation methods used, interpreted the data, and provided detailed interim and final reports needed to satisfy grant requirements. Involvement from the beginning of the project gave her valuable longitudinal perspective on development. She offered many suggestions that were implemented as KDN progressed. Had she been brought in at the conclusion of the project (after the grant money was spent), DCM would not have been as able to take advantage of her input.

Working with the KDN project manager, the independent evaluator developed a thorough evaluation plan very early in the KDN development process. The plan detailed her role in the process, the activities she would conduct, the data collection methods to be used, and a topical framework outlining the questions her evaluation would address. The topical framework was used to focus and guide the evaluator's observations and conversations with respondents, but as important new issues emerged, these were often explored as well.

The independent evaluator and DCM staff employed a naturalistic methodology to collect and analyze data (Lincoln and Guba, 1985). Naturalistic inquiry is particularly useful when seeking to understand real life settings and the situations and ways in which people experience them. It begins with a focus of study, and thorough systematic collection and analysis of the data identifies and reports on issues as they emerge. Data is collected from a variety of sources. This "triangulation" of data yields the completest possible description of the setting and program. Naturalistic inquiry often includes the respondents' own language as a way of painting a rich, detailed, and meaningful picture of their experience.

Ultimately, evaluation became a very collaborative and continuous effort. Although the independent evaluator was not present at all tests, she was present for most of the tests where major revisions to the program were assessed. She was also present at the meetings where the team took stock of the program's status and charted the course for the next stage of development and evaluation. The project manager kept her constantly informed via e-mail and phone. Information and opinions freely flowed to and from the evaluator, project manager, and Web designers.

The Evaluation Process

Although DCM already had an e-mail based math program on the Web and was experienced in conducting engineering programs for children, KDN was a venture into uncharted territory. Because there were so many aspects of the program to evaluate, it was helpful to isolate components and design evaluation to answer specific questions. Most components were assessed individually first and then incorporated into a trial of the whole KDN experience. For example, when the Slowly Moving Marble Challenge was being developed, DCM employees' children in the museum on a school holiday were the first to try their hands at designing and building a pathway to slow down a rolling marble. These children didn't use the computer; they simply built their pathways. A few months later, a 3rd-grade school class (taught by an KDN Advisory Board member) took on the marble challenge as part of the total KDN experience. Web site components were also evaluated in isolation first. Numerous versions of the Engineer chatboard were tested by the Web designers, then DCM staff, before children ever encountered them. The first children to see the chatboard were the project manager's children visiting the museum. Many iterations later, the first classroom trials of the chatboard occurred.

At the outset, three overriding questions were identified. Evaluation would be geared toward answering these questions:

1. Is KDN technologically feasible?

2. Would teachers and kids be interested in talking to the engineer?

3. What, exactly, should the content of the Web site be?

These questions and how we answered them are discussed below. Methods are summarized in Table 2.

Evaluation Techniques Used	Notes
Participant observation	Primary method of data collection; yielded the most valuable information
Teacher focus groups	Replaced teacher written surveys; teachers were interviewed in groups
Depth interviews	A few teachers and website programmer interviewed primarily by independent evaluator
Teacher usability studies	Observed teachers navigating website for the firs time; very valuable for identifying gaps in information and navigation problems
Website developer usability studies	Provided concrete suggestions on ways to improve site navigation and user-friendliness
Team meetings	Involved evaluator, DCM staff and web developer; critical for planning next phase of project development
Review of program documents	Project manager shared all documents with independent evaluator, including children's work and comments, teacher comments, conversations with other DCM staff
Written surveys	Used only at the beginning of project; poor return rate and few comments made them inadequate for program needs

Table 2: Evaluation methods used during KDN development.

Is KDN technologically feasible?

The first question to be answered was whether live two-way communication was even technically possible *in the school setting*. Three primary concerns were quickly identified:

1. The cognitive development of young children and their resultant ability to communicate would be key factors in determining the appropriate age to use KDN. KDN was initially aimed at children in 1st through 5th grades. The younger children in particular typically have poor typing skills, and have trouble communicating with someone not in the same room with them. The appropriate age for children to use KDN became one of the most significant questions to be answered.

2. Elementary schools typically do not have the latest in computer technology. It is not uncommon for computers to be more than 5 years old. Would KDN work on Windows 95 and Internet Explorer 3.5? At the time of development, many schools didn't have Flash technology. What could be used to make the Web site look appealing to kids and teachers? A key obstacle was firewalls in schools that block the streaming data necessary for the two-way communication at the heart of KDN. Would it be possible to tunnel through the firewalls, or would KDN use be restricted to school districts without them?

3. If the first two issues could be resolved, the safety of the children and consequent security of the communication became major issues. The Museum and KDN users needed assurance that children would be protected from inappropriate communications and unauthorized users.

Would kids and teachers be interested in talking to the engineer?

Assuming the technology issues were worked out, we still needed to know if students and teachers would find talking to the engineer on-line worth their time and effort. "Ask the expert" Web sites other than KDN are based on e-mail communication where children have plenty of time to formulate questions which are reviewed by teachers or parents who correct grammar and edit to ensure clarity. Would children be able to communicate well enough to gain usable advice about their gadgets? Finding out if students thought talking to the engineer helped them meet the challenge was a major focus of the first two field studies and a lesser focus in following studies.

Because children's access to KDN, and especially the engineer, depended on the cooperation of teachers and parents, the grownups' perception of the value of talking to the engineer was important to assess. Would teachers see enough value that they would be willing to negotiate the potential maze of scheduling computer time and learning to use the Web site?

What, exactly, should the content of the Web site be?

The initial Web site concept included the challenge presented in short story format, a printable student journal, a student-engineer chatboard with only text chat, and an on-line bulletin board where pictures of gadgets could be posted. The inadequacy of these features became apparent very early in development. Identifying what was necessary for effective student-engineer communication was a focus early in the project. Programming and testing the solutions ultimately took a year longer and $70,000 more than was originally planned.

Phase I: Field Studies

Evaluation techniques originally planned (based on the actual grant application)were:

- Review of materials by DCM staff;
- Pilot tests (field studies) with small groups of students in schools the first year;
- Field studies in full classrooms the second year;
- Written questionnaires and in-depth interviews for teachers with students participating in the pilot tests;
- Observation of pilot studies by DCM staff and the independent evaluator, and discussion with participants;

- Written parent surveys;

- In-depth case studies of selected students;

- Collection and evaluation of student materials, including journals, drawings, and video, to asses changes in the students' creative problem solving skills, ability to transfer knowledge to new situations, and ability to work as teams;

- Presentation of ongoing evaluation results to the Advisory Board and consultation with the members.

During the first year of the project, two small field studies were conducted at schools with which DCM had a strong relationship. The purpose was to assess students' reaction to the concept of an on-line engineer, the necessary components for effective communication, the appropriate age for KDN, and the appropriateness of the first challenge. Teachers chose a few students to participate in KDN. The museum brought in their own computers and set up a closed link between the engineer's computer and students' computer. Communication was via a FileMakerPro database made to resemble a Web page with two-way text chat in the first field study. The second study utilized a CU-SeeMe program to provide video chat with an engineer via a local area network

Much was learned about evaluation during these studies. Written surveys, so frequently used in evaluation, yielded minimal valuable information. Parents and teachers did not take time to record the detailed observations that are needed for comprehensive assessment. The value was also limited because the people completing the surveys (parents and teachers) were not present when children were working with KDN. Direct observation of students and conversations with them yielded the most information. Student journals were collected and an extensive video record of the sessions was made. Video and journals proved less valuable for evaluation purposes, but very valuable for bringing a new project manager up to speed a few months later.

A third small group trial aimed predominantly at testing a new challenge was held at the museum. No engineer was involved in this trial. A Girl Scout leader was asked to gather a group of students for this test. Children were asked to read two versions of the Puppet Challenge aimed at different age groups. The children were asked questions to see if they understood the story and its vocabulary, and if they found it interesting. The team wanted to assess the story format, the particular versions of this story, whether the challenge provided an appropriately open-ended engineering problem solving experience, and to revisit the question of the appropriate age for doing KDN. The session was videotaped and the children's journals were collected. The independent evaluator was not present. Once again, personal conversations with and observations of the children provided the most useful data.

These very preliminary field studies answered some basic questions and posed new ones:

- Kids liked talking to the engineer. They found the conversation helpful.

- Kids in 1st and 2nd grade lacked the skills necessary to communicate solely through typing. Third graders were marginal. The search for a visual and possible audio component to communication was on.

- It was going to take some fancy programming to pull this off on the Web.

- The way the challenge was presented needed work in order to make it clear and understandable. At that time, challenges were presented in a short story format involving a girl and boy and their teacher, Ms. Curious.

Results of the field studies were presented to the Advisory Board. They felt that the potential impact of KDN on science and engineering education was huge, and that there were still a lot of questions to be answered. Extensive discussions occurred about the best form of communication best (email, video, live text chat, voice, fax, or phone), the focus of KDN (science, computers, or engineering), and the importance of addressing teachers' needs.

A big meeting was held that included the newly hired Web site designers (Eduweb), the independent evaluator, the old and new project managers, and most of the DCM staff who had worked on KDN to date. A free flowing exchange of ideas made sure everyone was on the same page regarding where the project had been and what needed to happen in the near future. Eduweb presented information of the most appropriate technology available. An understanding of two overarching types of problems emerged from the meeting: technology problems and people problems.

Eduweb focused on finding technology that could provide the best communication, could tunnel through school firewalls, and was compatible with the typical aging computer systems common in elementary schools. Making the site user-friendly – in both graphic design and navigation - and compatible with school computers also became a focus. DCM would review everything and try out all interactive components of the site.

DCM focused on the people problems. Most of the team had already concluded that 1st and 2nd graders were too young to use KDN effectively. However, the Museum felt that additional data was needed to support this conclusion for the Illinois State Board of Education (the funder). It was important to make the teacher, parent and child user interested in and comfortable with KDN. This would be addressed through both the content of the Web site, and Eduweb's work on the graphic design and navigation. Finding and training on-line engineers continued to be of concern. DCM was also responsible for coordinating with schools which had volunteered to participate in the second round of field studies.

Over the next few months the first version of the KDN Web site was developed in preparation for the 1st full-fledged classroom tests. During this period, the DCM project manager served as the internal evaluator. Eduweb posted pages on the Web for the project manager to review. She typically reviewed the first few iterations of a Web page or design element, and then asked for opinions from other DCM staff. Preliminary assessment of firewall tunneling technology was tested on a few local school computers before being incorporated into the site. Any DCM staff children visiting the museum were pressed into service as voluntary evaluators. They tried out the functionality of the Engineer chatboard and drawing program, critiqued the graphics, and were usually the first to try their hands at a new challenge.

Because the site was to be accessed by entire classes of children at once, DCM staff participated in group tests of the drawing program and the Engineer chatboard. We found again and again that even if it works on the programmer's computer and even the project manager's computer, that does not mean that 10 or 20 or 30 computers could draw or use the chatboard at once. The programmer would access error logs and make adjustments, and we'd try again another day.

Phase 2: Classroom Tests

Scheduling classroom tests proved to be complicated; school calendars and schedules are generally very rigid. An extraordinary amount of accommodation was necessary on the part of the school. The whole school was often affected because computer lab time and frequently library time were adjusted so that the classes using KDN had access to the computers.

What we did right in these trials

Administrators at the participating schools were on the KDN Advisory board, so they were familiar with the concept behind KDN. These schools and administrators had, for the most part, been selected because DCM conducted many programs there and had a good relationship with the teachers and administrators.

Students from 1st through 4th grade were involved in this phase of testing, allowing collection of data on the age issue. Also important were the responses of teachers and students of different ages to the challenges and their presentation, the KDN student journals, and students' desire, ability, and willingness to talk to the engineer.

To address 1st and 2nd graders' communication difficulties, a drawing program and whiteboard with chat component (Engineer chatboard) had been added to the site to provide a visual component. Children and engineers could both see the child's design and mark on it. Cell phones were used with and without the whiteboard to talk to the engineer as the children were building.

Three very different schools were chosen: a suburban school in a middleclass neighborhood (School

1), a rural, smalltown school (School 2), and a suburban school with a large Spanish-speaking population in a low income neighborhood (School 3).

DCM provided a Spanish-speaking facilitator for the bilingual classes at School 3. This allowed us to ensure students understood KDN as well as allowing DCM to receive feedback from the students.

Although facilitation was by experienced DCM personnel, classroom teachers were present for all KDN sessions and actively participated in the KDN experience. DCM and the independent evaluator chatted with teachers before and during KDN sessions to find out what teachers wanted and needed before using KDN. Informal conversations allowed us to see how revolutionary open-ended, child-centered problem solving is to some teachers and what type of background information on the philosophy behind KDN would be needed on the Web site. For example, one 2nd grader making a bell puppet asked a DCM facilitator for help early in the process, and was later helped by her teacher. At the end of the session, the DCM facilitator asked her how her bell turned out. The girl said, "It's not a bell anymore. My teacher changed it." It was clear the Web site needed specific guidance on how teachers could help *children* solve their problems instead of taking over their project.

Eduweb's project manager and the computer programmer were present at the first school where testing occurred. Although both are experienced specialists in children's Web sites, they agreed there was no substitute for seeing children actually trying to use the site *en masse* in the school. It gave them a new appreciation for the state of technology in the elementary school, and of the problems of Web site functionality with young children.

Eduweb's computer programmer and project manager both played the role of engineer several times during this phase of testing. This gave them first hand knowledge of the difficulties of using the engineer's communication tools. The programmer reported that this experience was extremely valuable. For example, being limited to only the student whiteboard chat for communication, the engineer didn't know what was happening in the classroom during a long period of inactivity. Had the students all left for lunch? The solution was to add a separate chat component enabling a private channel of communication with the teacher.

Major revisions to the drawing program and the Engineer chatboard were made after each school participated.

The independent evaluator was present for most of the session at all three schools. She carefully observed all aspects of the sessions, chatted with students and teachers during the sessions, and asked follow-up questions at the end. For those sessions not attended, the project manager provided detailed notes and comments. Following the evaluator's model, the project manager asked students and teachers follow-up questions at the end of their KDN experience.

After School 1 tests, the team completely revised the challenge format to make it clearer. After in-house evaluation at DCM, the new format was tested at School 2, revised slightly based on teacher comments, and tested again at School 3. This format is still used today.

What we did wrong in these trials

The schools had been chosen because their administrators were on the KDN Advisory Board, but without regard for the state of their computer systems. A few weeks before the first tests, we found out School 1 had very old computers that simply could not handle even the most basic aspects of the KDN site. Eventually, the district computer expert borrowed a portable iBook lab from the junior high and set it up at the elementary school for the KDN tests. Consequently, the children were learning to use a new computer and the KDN Web site at the same time.

Because the administrators, rather than the teachers, had made the decision to participate in KDN development, many teachers were less than enthusiastic. Many, but not all, were not adept at computer use and did not see the value of hands-on activities. However, one bilingual teacher at School 3 enthusiastically reported that KDN had sparked a whole class discussion on how "this concept of keep trying until you get it right" applies to the rest of life, too. He said that the open-ended problem solving of KDN helps teachers see a new way of doing things, too.

School 3 had no computer lab at all until two months before the KDN trial. Most of these children had no home computers and consequently very little computer expertise. This group of students, however, proved to be very excited about KDN and vividly illustrated the potential impact of KDN on traditionally underserved populations. The students thrived when building their gadgets because there was no language issue with the hands-on activity. The students easily learned to use the Web site's interactive features. They were thrilled to talk to the engineer, and the girls especially were excited to learn that the engineer was a woman. However, many of their teachers struggled with the computer. These teachers, many of whom had never used computers in their classrooms, illustrated the importance of thorough, clear guidance in using KDN and user-friendly navigation.

Because of individual school scheduling constraints, the first trials were with 1st and 2nd graders, the group we expected to have the most trouble with KDN in general. A more refined version of the Web site would have given us a truer picture of the children's abilities without the complications of computer bugs. However, the tests still demonstrated conclusively that these young children were not developmentally ready for KDN.

Follow-up questionnaires were given to the teachers. Due to the poor return rate on written teacher questionnaires, the evaluator scheduled group debriefing sessions. This was a more efficient use of teachers' time, and the group dynamic yielded more useful information. For example, a teacher would raise an issue, and several others would agree. This helped us see which issues resonated most with teachers. These discussions provided valuable, specific data that was used by the project manager and programmers to make major changes to the Web site and program content.

Web site Usability Tests

After the classroom tests at the third school, the KDN team met and concluded we basically had a workable Website, albeit still with many bugs. A primary concern was that the site had become so big that navigation was a mess. The home page was overloaded with information, and there was a general lack of user-friendliness to the site, largely due to its complexity. Because the team was so accustomed to using the site, a new pair of eyes was needed for assessment. Two new evaluation strategies were adopted: teacher usability studies and review by independent Web site developers.

Teacher Usability Studies

To identify gaps in information and get teachers' perspectives, two elementary classroom teachers were brought in to look at the Web site in two usability studies. We gave the teachers very little orientation to the site because we wanted to observe them the first time they saw the Web site. Using a "think-aloud" protocol, the evaluator, project manager and Web developer took careful notes as each teacher navigated the site while talking aloud. This gave us specific information on where navigation and information needed revising. Following their time exploring the site, we held debriefing sessions during which teachers explained and expanded upon some of their comments. These usability studies were among the most useful techniques used in the entire evaluation.

After the first teacher reviewed the site, major revisions were made. Then a second teacher was observed using the same "think-aloud" protocol. After a morning of learning about the site, we observed as the teacher directed six children, who were unfamiliar with KDN, in using the site, including the drawing program and talking to the engineer. As the focus of this exercise was on the Web site itself, the children did not construct their gadgets. One of the most striking observations was that, prior to observing the children, the teacher did not think they would be interested in talking to the engineer. She stated she was very surprised to see how enthusiastic they were when talking to the engineer, and this changed her opinion of KDN considerably.

Both teachers pointed up a paradox the KDN team had long been wrestling with: they thought there was too much information on the site, but they wanted more information. The KDN team discussed this problem and concluded the solution must lie in presenting the information in a more eye-appealing way. Once we had a good idea what teachers needed on the site, we needed an uninitiated Web designer's opinion on how to give it to them.

Independent Web site Developer Reviews

After extensive revisions based on the teacher usability studies, two Web site developers unfamiliar with KDN were asked to review the site. We were particularly interested in their ideas on the overall look and feel of the site – especially graphic design and navigation. Although there were some common themes in the reviewers' comments – especially the need to simplify the navigation and put less information on each page – the reviewers came up with significantly different solutions to these problems. The KDN team evaluated the reviewers' comments and used them as a starting point for major revisions to the site navigation, page layout, and site organization.

Analysis and Redesign

Much discussion occurred between Eduweb and DCM, and many prototypes were reviewed before a basic design evolved. We were, in effect, reinventing the wheel with four new visions on how to go about it – DCM's, Eduweb's, and the two reviewers'. Text on every page was revised, a new graphics shell was developed, site content was reorganized and an entirely new navigation scheme developed, colors were constantly changed, a new log-in sequence was developed, and major revisions - both functional and cosmetic - to the drawing program and engineer chatboard were implemented. Revising the KDN home page was a major effort. The home page used in Phase 2 classroom trials is shown in Figure 2, and the current home page is shown in Figure 3. There were at least a dozen interim versions. The resulting current version of the site bears little resemblance to the reviewers' suggestions or the version they reviewed, but based on feedback from teachers who have since used the site, most of the problems have been resolved.

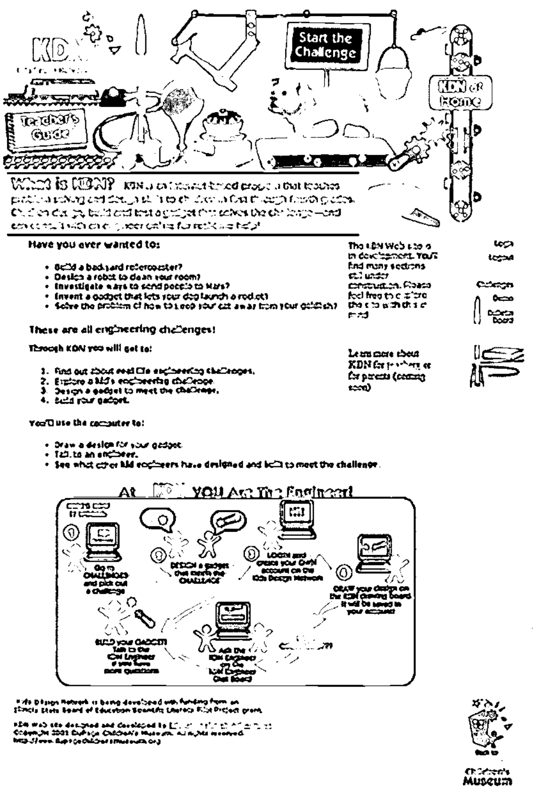

Figure 2: KDN home page used in Phase 2 classroom trials.

This page was too long, too busy and generally notuser-friendly.

The solution to too much information was the creation of a KDN "Quick Tour" that takes visitors through a virtual tour explaining how KDN works; and a "What is KDN?" section providing additional details.

Phase 3 Classroom Trials

The plan for Phase 3 Classroom Trials was for class-

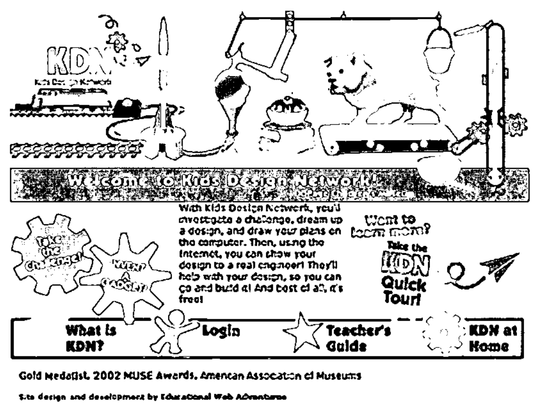

Figure 3: The current KDN home page (http://www.dupagechildrensmuseum.org).

room teachers to facilitate KDN with only e-mail and phone support from DCM. Two teachers, one in a small town near Chicago and one in an inner city Chicago school, found KDN on the Web on their own. They provided the first real test of the extensive information in the Teacher's Guide section of the Web site, and the newly revised Web site. The schools' proximity to DCM allowed the KDN Project Manager and, in one case, the independent evaluator to visit the schools and observe as students built their gadgets and talked to the engineer. The evaluator and project manager conducted an extensive interview with the teachers. Their feedback was used to further refine the Web site.

We were relieved to see that these children's experience was remarkably similar to KDN activities facilitated by DCM staff. The teachers were able to glean the philosophy behind KDN from the Website, navigate the site, and use its information to produce a hands-on, minds-on learning activity for their classes. Each teacher had added an individual twist to the challenge to emphasize different concepts, but the child-centered, creativity and problem solving concepts were paramount. The teachers had successfully used KDN exactly as was intended.

However, there were still problems with the engineer chatboard, most notably stemming from a modification that allowed the engineer to simultaneously chat with several students. Ultimately, this feature was disabled because it did not function reliably.

Conclusion

Although the basic concept behind KDN has remained constant, KDN as originally envisioned by DCM was very different from today's version. From the early stages of development, the team placed emphasis on continuous evaluation and a willingness to make revisions. The collaboration in evaluation between professional evaluator, Web site developers, museum personnel and engineers resulted in a fluid development process that mimicked the engineering problem solving process KDN is designed to teach. This iterative nature of engineering design as applied to KDN development ultimately produced a user-friendly experience consistent with the Museum's mission. Our efforts were rewarded with a Gold Muse for Programs Emphasizing Two-way Communication from our peers at the American Association of Museums. However, the best rewards come from the children who say "WOW! This is so cool! I wanna be an engineer when I grow up!"

References

Lincoln, Y.S. and E.G. Yuba, 1985. *Naturalistic Inquiry.* Sage Publications, Beverly Hills, CA.

Virtual Visiting

The SEE Experience: Edutainment in 3D Virtual Worlds.

Nicoletta Di Blas, Politecnico di Milano, Italy; Susan Hazan, The Israel Museum, Jerusalem, Israel; and Paolo Paolini, Politecnico di Milano, Italy

Abstract

Shared virtual worlds are innovative applications where several users, represented by Avatars, simultaneously access via Internet a 3D space. Users co-operate through interaction with the environment and with each other, manipulating objects and chatting as they go.Apart from in the well-documented online action games industry, now often played out in real-time over the Internet, the success stories of virtual worlds and shared VRML environments have been sporadic. Interest wanes from an enthusiastic beginning: the 3D world looks both enticing and promising, but the attention curve drops sharply as the users find the space either too difficult to navigate or not satisfyingly engaging.The key to the success or failure of these environments relates directly to the content offered and the ability of the environment to *hold* a critical mass of users in real time.This is the desperately sought after magic social glue that determines the experience to be meaningful and actively keeps the users under its virtual spell.

The paper will present an innovative project, a shared virtual world geared to schools in an educational and highly structured,environment. SEE, Shrine Educational Experience, is a co-operative project developed jointly by Shrine of the Book at the Israel Museum, Jerusalem, and the Politecnico di Milano, Italy, offering a pioneering educational environment based on a shared 3D virtual world, where classes from all over the world meet in order to learn and discuss issues related to the famous Dead Sea Scrolls and the tiny community (probably of Essenes) who once lived by the Dead Sea.

Keywords: e-learning, 3D virtual worlds, edutainment

1. 3D Worlds: A New Challenge to Museum Communication

Exactly two years ago, at Museums and the Web 2001 in Seattle, Susan Hazan of the Israel Museum, Jerusalem, and Professor Paolo Paolini and his team sat down for a micro-brain-storming session to think of a way to harness the impressive technology developed for their "Leonardo da Vinci" project (Barbieri, T., P. Paolini et al,. 1999) for the Israel Museum in Jerusalem.The teams soon focused on the idea of a real-time, educational scenario that would extend the museum's mission to educate and interpret collections for its online community, and the seed was sown.

The Education Department of the Israel Museum, Jerusalem, is proud of its track record of innovative, online educational activities. In 1996, during the exhibition, *Children of the World Draw Jerusalem at 3000*, the department developed an ambitious project to provide a unique opportunity for children from all parts of the globe to take part in a drawing contest on the theme of Jerusalem (Hazan, 1996). For the duration of the exhibition, the department set up three Macintosh computers. For six weeks they were staffed by high school students from our Multimedia Education Unit with a specially designed Web site that offered remote visitors the opportunity to upload their own artworks, to be displayed online, as well as to play an interactive trivia game to learn about Jerusalem, its history, geography and religions. A further opportunity to 'visit' the exhibition was made possible through a video window. Using the software CUSee-Me, the exhibition, with local visitors as spontaneous players on the live stage, was broadcast from the gallery for 8 hours a day for the full six weeks of the exhibitions, at the same time receiving remote visitors to the museum in real-time.

The exhibition *In the Light of the Menorah: Story of a Symbol*, was launched in 1998, a photo-realistic 3D gallery tour streamed at 15 frames per second over standard dial-up modems with embedded hyperlinks enabling students to interact with the movie (Hazan, 1999). This was an adaptation of a major temporary exhibition held at the Israel Museum, one which traced the manifold incarnations and interpretations of the seven-branched candelabrum from Biblical sacred object to national emblem, as represented in objects from the Museum's extensive collections of archeology, Judaica, and the fine arts. In the same year the Israel Museum forged a partnership with the Hebrew University and a number of Israeli cultural institutions in order to develop a specially devised Web site, *Galim,* meaning waves. Via the Museum@school project ,*Galim* students navigated the 3D environment through online games, quizzes and self-directed exploration, delivering a new museum experience straight into the classroom.

The museum was confident that if these two components, the real-time excitement of the video window and the challenge of wandering a 3D gallery online searching for clues, could be combined, the result could be a novel way of extending the museum's mandate to enhance and interpret its collections for remote visitors. The magic created by the Politecnico world for the Leonardo project appeared to present exciting new possibilities.

Games too turned out to be a powerful source of inspiration. *MechAssault,* billed as *fierce tactical combat and destruction, 31st-century style,* is now *Live*-enabled, meaning that the real-time component allows users to play with each other online as many of the action games already do, and even if youngsters are not out there shooting each other, or forming online strategies to blast each other out of the game, many of the online activities are just as superficial. The popular *The Sims Deluxe* (http://thesims.ea.com/), the latest and wildly popular version of the original *The Sims*™, is promoted as *Create an entire neighborhood of Sims and run or ruin their lives,* while NFL Fever, the best-selling football game on X-Box, boasts *some of the most bone-crunching, realistic animation of any football video game to date.* Many less virulent environments have been developed using VRML. Clearly "*being there*" in a real-time interaction with other users does hold a potent attraction, but performing specific activities in shared 3D spaces has proved to be difficult and sometimes frustrating. Once users have negotiated the technological obstacles to enter the online environment, (with a fairly robust computer and speedy communication), the key to the success or failure of these environments directly relates to the content offered and the ability of the environment to *hold* a critical mass of users in the same space at the same time. This is the desperately sought after magic glue that makes the experience meaningful and actively keeps the users under its (virtual) spell. Clearly the museum is not considered a player in this ballpark – but on the other hand, should the museum even want to be there? What can a content rich provider such as the museum, with an impressive track record for 'holding' its visitors in compelling educational scenarios, possibly offer this industry?

The Israel Museum made a decisive move into this new arena, selecting the most valuable and compelling of its wide ranging collections as the thematic core of the activity – the world renowned Dead Sea Scrolls. The Dead Sea Scrolls are on permanent display at the Israel Museum, Jerusalem, in a stunning building designed by the American architects Armand Bartos and Frederick Kiesler. Dr Adolfo Roitman, Head of the Shrine of the Book and Curator of the Dead Sea Scrolls, brought together a team of museum educators and curators, as well as staff of the New Media Department, to explore the curatorial content of the SEE project developed from his museum expertise as author of the exhibition, *A Day at Qumran: The Dead Sea Sect and Its Scrolls* (Roitman, 1997) and in his role as gatekeeper of the precious collections. The Shrine of the Book is located at the heart of the political and cultural center of the State of Israel, close to the *Knesset,* (Parliament), government offices, the Hebrew University, and the Supreme Court Building. A white, ceramic covered dome covers the galleries, which are two-thirds below the ground and house the permanent collection. The striking juxtaposition of the gentle curve of the dome and the black, angular, basalt wall, both stark yet opposing geometrical shapes, creates a stunning vista and contributes a unique architectural phenomenon to the Jerusalem horizon. This was the portal selected for the online virtual world, a space for the social and educational interaction to be staged. The magic of the virtual world was that other worlds could be brought into the constellation of this experience ,and so the ambience of Qumran could be brought into the historical narrative and into a new reality.

The curatorial and educational team of the Shrine of the Book at the Israel Museum, Jerusalem, is con-

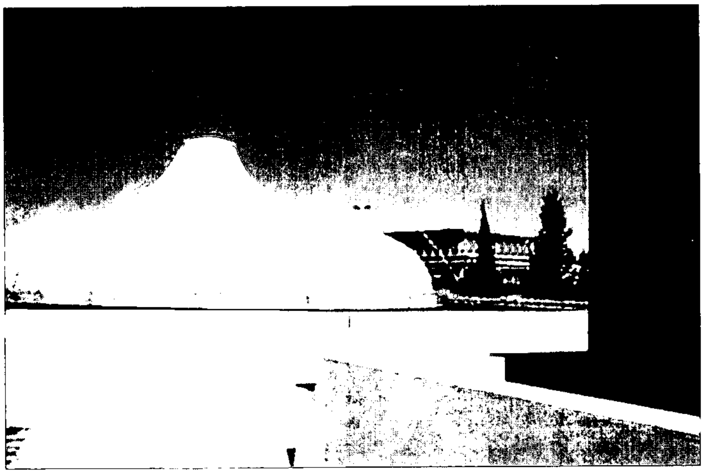

Fig. 1:The Shrine of the Book with the Knesset (Parliament)

cerned with preserving not only the material evidence of the scrolls themselves but also, just as important, the message contained within them.The SEE project, perhaps one of the first museum virtual worlds to be developed, is an exciting new and dynamic environment where objects can be manipulated and hidden portals accessed while avatars move around the shared space while socializing with each other in a real-time.

2. The *Virtual Leonardo* Project: Lessons Learned

Virtual Leonardo was deployed for the first time in 1999 (Barbieri, Paolini et al ii, 1999; Barbieri & Paolini, 2000} with a new version launched in 2001; it was developed by Politecnico di Milano for the Museum of Science and Technology, Leonardo da Vinci, in Milan.

The original objectives and the main features of the project are perceived as follows:

A 3D graphic environment to be used as a "container" to display "objects" of the museum.

Specific choices in our case were an "idealized" rendering of the cloisters of the museum (hosted in a former monastery) as the container, and the "Leonardo machines" the objects.The latter are 3D interactive objects, representing actual wood machines on display in the museum; the machines were built several decades ago in order to "render" ideas and concepts developed by Leonardo in several of his drawings.

Users are encouraged to visit the "virtual museum" in groups

The users are encouraged to visit the virtual museum in groups: a family, a group of friends, a school, and few visitors with or without an expert. This creates a sense of "not being alone" while visiting the museum, a factor that increases fun, engagement, and interaction, while making the overall visit more effective and rewarding.

Avatars represent users who are visiting the "virtual museum"

Users "visiting" the museum are actualized through avatars, i.e. graphic objects that appear as either real persons, or graphic renderings.The purpose of using avatars is twofold: it allows the users to perceive their own physical positions in the virtual space as well as to recognize other avatars' positions within the museum. Users move their Avatars in order to change their position within the museum or to make contact with the objects.

Advanced functions allow users to cooperate in several fashions

We realized, since the inception of the project, that providing "traditional" interaction mechanisms within the virtual museum was not sufficient.Traditional 3D environments allow users to chat, to interact with objects, and to wander around (using "real world" limitations, in general). We observed that in most 3D environments the above possibilities did not encourage a great deal of "cooperation" among users: chat was the main instrument of cooperation with the others, making avatars almost useless (i.e. not exploiting the spatial cooperation). We strongly felt this concern: how could we get the visitors to actually interact with each other, besides chatting? We therefore introduced the notions of "cooperation capabilities" and "cooperation metaphors". A cooperation capability defines something that users can do: e.g. move, speak to everybody, whisper to a few, change visualization point of view (using fixed cameras, for example), make hyper jumps, look through the eyes of someone else, etc. A cooperation metaphor (Barbieri & Paolini, 2001 – the term "metaphor" is not really appropriate, but this was the term that we historically used)) is a combination of cooperation capabilities assigned to the different users, in order to obtain a specific overall behavior: e.g. during an "explanation" the "guide" can "speak" to all the visitors, while each visitor can speak only to the guide and not to the other visitors, and hyper jumps are not allowed. Cooperation metaphors were understood through the experience of *Virtual Leonardo*, but not fully exploited within it (more advanced use of cooperation metaphors has been introduced in later applications developed by Politecnico).

A "museum guide" takes the visitors around

We did realize very quickly that a group of "visitors" in a virtual museum, without a specific goal in mind or a task to perform, very likely would "hang around", without actually knowing what to do. Therefore we introduced the notion of "museum guide", with a number of duties:

- stimulating the curiosity of visitors
- organizing the visit
- providing additional information
- answering questions

If audio guides and "physical" guides are a popular option in a real museum, they are indispensable in a virtual museum for a number of reasons that are probably obvious to all virtual visitors who have experimented the feeling, "what am I going to do next?" and even more obvious when a group of visitors must be kept together for a meaningful visit.

Visiting a "virtual museum" becomes a "social" activity

Thinking about all these factors together and understanding that visiting a real museum alone is a different experience from visiting with other people (with family, friends, experts), made these lessons insightful in our future development of shared virtual worlds.

Virtual Leonardo has raised the interest of practitioners, researchers and the general public. Its simple idea is appealing to all: "*at this museum you can bring a date*", proclaimed the title of an article on New York Times online (Mirapaul 1999). So, the idea of visiting a virtual museum evolved rom a "one at a time" activity to a social activity. Many users have visited *virtual Leonardo*, but we realized that there were some flaws regarding the three main aspects of the application: the interaction, the content, and the reproduction of the real world limitations. More specifically, we noted the following problems:

About the interaction:

- Visitors mainly interacted with the guides and the objects; there was little reason for a visitor to interact with another visitor.
- There was an inherent tradeoff between having the visitors spend their time interacting with the world and the objects or having them interact with each other.
- Despite the power of cooperation metaphors, it was still difficult (except in demos) to avoid avatars hanging around, while the focus of the interaction was the chat.

About the content:

- The 3D environment was not appropriate to deliver either large quantities of information or high quality visual information.

- If the main emphasis of the visit were to acquire new information, there would be limited opportunities to get to know other visitors.

- Whenever the visit was geared to emphasize discussion among visitors, clearly users should be supplied with enough material beforehand to be able to engage with the material during the session.

About the reproduction of the real world constraints:

- Mimicking the physical limitations of the real world turned out not to be useful in general, and quite boring in the end.

If the virtual visit to a museum is to provide a meaningful and forceful experience, something more structured and compelling is needed. Therefore, when the Israel Museum asked Politecnico to start to investigate a new project targeting high and junior high school pupils and focused around the Dead Sea Scrolls and Qumran, the Politecnico staff decided to create a completely new kind of experience built on the foundations of the *Virtual Leonardo* project.

3.The Shrine Educational Experience

For the new application, baptized "SEE" (Shrine Educational Experience), we chose high school and junior high school students as our target group. In the future, we are considering other target groups: such as tourists who are interested in the Holy Land, scholars who want to *meet* in the virtual world to discuss an issue related to Qumran and the Dead Sea Scrolls, or simply well educated *surfers*. The SEE experience is based on the interaction of four classes, from different parts of the world, virtually *meeting* in real-time in a 3D virtual environment, together with the museum guide. The 3D world consists of several different settings, including (Fig. 2) the internal and external (real) spaces of the Israel Museum's Shrine of the Book where the Scrolls are on display, while reiterating the potent symbolic metaphors of the buildings themselves (Fig. 3) set in contemporary Jerusalem, with a hyperlink to the Qumran desert, and the historical setting of the Essene community.

The students are represented by avatars and move, play and interact in the virtual world. Meetings last about an hour, where content (made available in advance to schools via a tutorial Web site) is recollected and applied in the "cultural games". SEE thus offers an *edutainment* experience, where students engage with the cultural content while enjoying themselves at the same time. The entire SEE experience consists of 3 interactive sessions, where students are led to discover respectively (1) the Shrine of the Book Campus in the Israel Museum in Jerusalem, (2) the Essene community who once lived in Qumran, and (3) the connections between the Qumran community with other religions and culture.

When SEE was designed, four major requirements were considered:

1. The relevance of the cultural content, the knowledge about Qumran, and related issues,(religious, historical, technical, social, etc.).

2. The inter-cultural exchange through *meetings* in the virtual space, with students of different countries and/or of different cultures.

3. The interaction between physically remote participants in an innovative and engaging experience, through the virtual environments. The games that

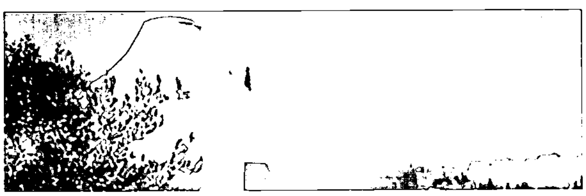

Fig. 2: The Israel Museum's Shrine of theBook

Fig. 3: A screenshot of the virtual Shrine external environment

the students are invited to perform in also consolidate "team-ship", while forging new relationships and ties among different schools.

4. An acquaintance with state of the art Information and Communication Technologies and the possibility of integrating modern multimedia, graphic, Web and Internet technologies into the school program.

The design improvements were adapted while accommodating the above self-subscribed goals as well as through what we had learned through the *Leonardo* experience and consequently integrated into the SEE application.

4. From *Leonardo* to SEE

In this session we go through the three "weak points" spotted in the Leonardo project; that is, how to promote the interaction, how to deliver the content, and what constraints of the real world should be reproduced, showing how we tried to overcome them in planning the SEE experience.

Promoting the interaction

One of the well-known drawbacks of 3D worlds is that visitors very soon get bored of moving around the environment when there is little motivation to do so. If their time is spent simply chatting with the others, this becomes an overall unsatisfying experience. Therefore we decided to carefully plan a set of activities through a storyboard to be acted out during the SEE experience, in order to avoid wasting precious online time or avatars hanging around aimlessly in corners.

Basically, each online session is divided into four main phases:

1. A brief moment of *Introduction*, during which students meet each other in the virtual environment and are welcomed by the Guide.

2. An *Exhibition* phase, in which either new content is delivered or already known material is recollected through boards (that is, pop-up browser windows showing multimedia content in a traditional, readable 2D format).

3. Two *Cultural Games* where students, divided into two teams, have to prove their ability in a multiple-choice quiz (one student per team must perform an ability-game in order to win the right of answering first for his team; fig. 4) as well as a Treasure Hunt (fig. 5) or matching-pairs game, both taking place within a specifically dedicated space, similar to a labyrinth.

4. Finally, the Guide gathers the students for a short *Wrap-up* where the key issues of the session are summed up, homework is assigned, and the next session's main theme is briefly anticipated.

Our preliminary tests in schools have demonstrated that the *discipline* of the storyboard has the positive effect of offering students a sequence of tasks to perform, alternating with games and brief lectures and the question-answering sessions with the guide. The whole experience therefore is about following a brief - and is competitively exciting.

While we succeeded with *Virtual Leonardo* when visitors interacted with the objects, we failed in that they did not interact with each other apart from

Fig 4: Screenshot of an avatar approaching the first ability-game in the Quiz Space

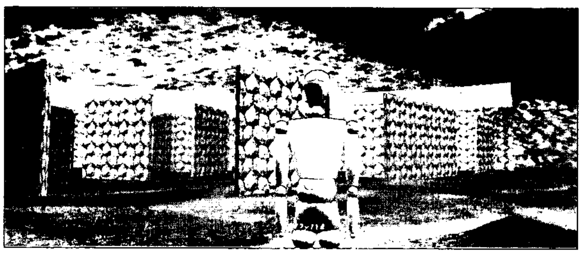

Fig 5: Screenshot of an avatar exploring the labyrinth during a Treasure Hunt

the chat. Therefore in planning the SEE experience, one of our most important issues was to compel avatars to interact (1) with the environment, (2) with the objects, as well as (3) with each other.

As regards points (1) and (2), all the games were planned so as to compel students to actively use the environment features, even the most advanced of the virtual world. Avatars may fly, see from another avatar's point of view (using his or her own eyes or another's), or view the site from fixed cameras. For example, in the Treasure Hunt students move around in the environment and find the sets of objects related to each other. In the brief time allotted to this challenge, looking through another's eyes and from another's point of view can be very useful, and even imperative, in order to be able to compare the objects to see whether they match or not before their time is up. In the quiz, students have to perform ability games (such as jumping from one platform to another without falling, pushing geometrical objects into holes), to win the right to answer to the questions.

In order to encourage avatars to collaborate with each other (point 3), to gain advantage in the games, students have to exchange opinions and really put their heads together. In the Treasure Hunt, for example, they have to check whether the object they've found matches the others, as all of the objects that the team needs to collect are semantically linked by a clue (i.e. *find all the objects that were used during the community's communal meal*). The quiz also encourages cooperation, for the right answer among three has to be spotted.

Tests in schools have shown that students love collaborating a team and interact not only for *strictly business* reasons as time is so short (find the correct answer, check the objects), but also to encourage their teammates with friendly messages to urge them on in the game.

Delivering cultural content

3D worlds are clearly not appropriate architectures for disseminating large amount of information, as neither lectures nor expansive documents work well with new means of communication. These limitations ironically encouraged creative cooperation and interaction among the visitors. Our challenge became therefore to match this simple observation with the educational goal of the application: how to deliver robust cultural content appropriate for a museum experience to students. Making background material available to schools through a *traditional* 2D site appeared to provide the solution, and the 2D site now functions as a gate to the 3D environment.

The background material provided consists of:

- Interviews with experts of the Dead Sea Scrolls and the Qumran community.
- Editorial insets explaining issues, events or characters in detail, mentioned by the experts in the interviews
- An anthology collecting all the excerpts from the Scrolls, the Bible, historical sources or other texts mentioned or quoted in the interviews
- Auxiliary didactic material providing background information on historical and geographical issues that may be obvious to some parts of the audience, and rather obscure to others (e.g. Israeli

Fig. 6: Screenshot of the Quiz Space, an environment dedicated to one of the Cultural Games

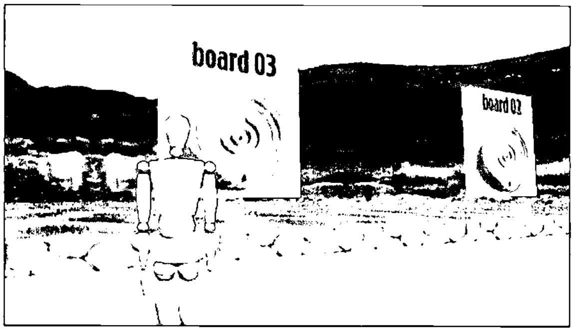

Fig. 7: The Guide's avatar near a board hotspot

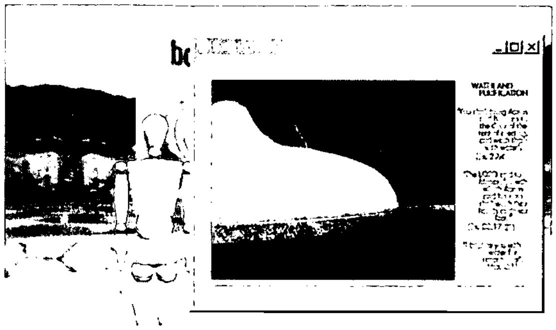

Fig. 8: Screenshot of a board

students know quite well where the Dead Sea is; for remote students this is not obvious at all).

Both the boards accessible though the 3D environment as well as the games themselves recall the relevant concepts and themes of the background material, summarizing a few key points and provoking online discussion with both the other participants and the guide.

The core of the background material consists of interviews with scholarly experts on the Dead Sea Scrolls and the Qumran community. Interviews are readily available both in full text and in summary. The interview format was selected for a number of reasons:

- The format is more immediate and straightforward, if compared to standard schoolbooks.
- Interviews provided a platform for different ideological, religious and cultural points of view. We have already interviewed (and plan to interview in the future) experts representing all different cultural and religious points of view on Qumran and related issues: Hebrew, Christians of all denominations, Arabs, lay scholars. Students are thus allowed to compare different approaches on the same topic.
- The *virtual* debate among scholars in the interviews enhances critical thinking. In fact, many issues concerning the Dead Sea Scrolls are currently under debate (for example: is the Q7 fragment an early testimonial of the Gospel?). Instead of arbitrarily presenting and promoting one opinion among the others, we decided to create a sort of "virtual panel" among experts on the hottest topics.
- Students may compare different opinions of a number of experts while realizing that many issues have not yet been resolved. In "normal" schoolbooks, students are led to think that all the issues are settled once and for all, whilst real historical and scientific research presents large fields still to be explored!
- Students may discover the typical dynamic of research. This is the result of the above points.

Reproducing the Real World

The observation that students tend to interact mainly via chat has recently brought us to some theoretic developments concerning what we called "proxemics semiotics"; i.e. the investigation of the "signals" that human beings exchange in the real world, often because of the physical limitations of the real world itself. "Proxemics" derives from the Greek verb *pros-semaino*; that is, "to mean together with", to mean *a coté* of the main communication channel, that is, the verbal channel. Proxemics semiotics is being developed by a joint effort of USI (University of Italian Switzerland) and Politecnico di Milano, with the goal of creating a new generation of more effective 3D virtual spaces (Di Blas & Paolini, 2003). For example, if we enter a room in a museum and see that all the visitors are gathered around a specific work of art, we are led to think that this must be an especially important work, possibly the most important in the room. In the virtual world, we can eliminate any physical limitation: but in doing so, we also cancel all those signals that in the real world are so meaningful! If avatars can look at a board from any part of the environment, all the others will never notice that they are interested in that particular board, unless motivated to do so by the others via chat. Otherwise, if the avatars were compelled to approach an object or a board or another avatar in order to interact ,this intention would be made clear immediately to all the others. This is one of the reasons that the use of chat during the experience is so thick: many of the things that in the real world are self-evident, because of the proxemics, ihave to be stated explicitly in the virtual world through a verbal message (that is, via chat). Avatars wanting to attract the attention of another must first say: *I'm looking at board 3, near the stairs* before declaring the reason of (*the meaning of the white dome is amazing!*).

We are currently considering which of the physical limitations of the real world should be maintained (or maybe somehow *translated*) in the 3D world, in order to enrich the overall communication and also to move it from the now almost exclusive chat-channel to the proxemics channel .We look forward to developing this exciting project in the near future.

5. Conclusions

SEE provides an educational experience centered on a Museum, with a number of distinctive features and innovative aspects:

- There is a blend of offline activities (background preparation, homework, etc.) and online activities.
- Online activities are a blend of "traditional" one-at-a-time navigation and cooperative activities.
- Cooperative activities emphasize the engagement of meeting other people and interacting with them. Information gathering is performed either offline (reading downloaded documents) or online.
- Background material is mainly based upon interviews, therefore purposely non systematic, possibly incomplete and inconsistent (sometimes experts have different opinions about interpretations or facts). The goal is to encourage students to understand the contradictions and the inconsistencies of the "real world" as opposed to the apparent and fictitious "cleanness" of textbooks.
- Via the background material and online experiences, students are led to correlate museum objects (i.e. the scrolls) and ancient culture to contemporary life, beliefs and experiences.
- The 3D environment is used with a well-defined storyboard (defining the sequence of scenes and of actions within the scenes), in order to keep the students "engaged" and under control.
- Competition (with the natural desire of winning) is used in order to encourage cultural preparation of the students.
- Gaming ability is also necessary to win; therefore students (rather than teachers) are encouraged to learn how to use advanced 3D-cooperation features.
- Last, but not least, the interaction among students of different cultures and backgrounds (with all the dangers of conflict always possible) is a major goal of the overall experience.

We are now working with very successful experimental usage (in Italy and Israel) and we will make SEE accessible world-wide in the fall of 2003. Only real-life usage will tell us whether our objectives were actually met or not.

References

Barbieri, T. (2000). Networked Virtual Environments for the Web: The WebTalk-I and WebTalk-II Architectures. *IEEE for Computer Multimedia & Expo 2000 (ICME), 2000*, New York, NY

Barbieri, T., F. Garzotto, et al. (2001). From Dust to StardDust: a Collaborative Virtual Computer Science Museum. *ICHIM 2001 International Cultural Heritage Informatics Meeting, 2001*, Milano, Italy

Barbieri, T. & P. Paolini (2001). Cooperation Metaphors for Virtual Museums. *Museums and the Web, 2001*, Seattle, WA.

Barbieri, T. & P. Paolini (2000). Cooperative Visits to WWW Museum Sites a Year Later: Evaluating the Effect. *Museums and the Web, 2000*, Minneapolis, MN.

Barbieri, T., P. Paolini et al. (1999). Visiting a Museum Together: how to share a visit to a virtual world. *Museums and the Web, 1999*, New Orleans, LA. 27-32.

Chang, N. (2002). The Roles of the Instructor in an Electronic Classroom. *Educational Multimedia and Hypermedia*, 2002, Association for the Advancement of Computing in Education, Denver, CO.

Di Blas, N. & P. Paolini (2003). Proxemics Semiotics in 3D worlds. (To be submitted to) *Studies in Communication Sciences*.

Falk, J.H., & L.D. Dierking (2000). *Learning from Museums. Visitor experiences and the Making of Meaning*. Walnut Creek, CA: Altamira Press.

Gardner, H. (1983). *Frames of Mind. The Theory of Multiple Intelligences*. New York, NY: Basic Books Inc.

Hazan, S. (1996). Children of the World Draw Jerusalem - Museum Video Window, MCN Spectra, A Quarterly Publication of the Museum Computer Network, Volume 24/2 winter. New York, NY, USA, 1996/97.

Hazan, S. (1996). *The Virtual Art Experience - Beyond the Museum Walls*, ICOM, CECA '96, New Strategies for Communication in Museums, Proceedings of ICOM/CECA '96, WUV, Universtatsverlag, Vienna, Austria.

Hazan, S, (1999) Linking and Thinking - The Museum@School, Selected Papers from an International Conference, Museums and the Web 1999 - New Orleans, Louisiana, March, 1999", Edited by David Bearman and Jennifer Trant, Archives & Museum Informatics, Pittsburgh, USA.

Mirapaul,. M. (June 1999). At this virtual museum you can bring a date. The New York Times On The Web.

Roitman, A. (1997). *A Day at Qumran: The Dead Sea Sect and Its Scrolls*. Jerusalem: The Israel Museum.

Make Your Museum Talk: Natural Language Interfaces for Cultural Institutions

Stefania Boiano (freelance), Giuliano Gaia (freelance), Morgana Caldarini (Jargon), Italy

Abstract

A museum usually talks to its audience through a variety of channels, such as Web sites, help desks, human guides, brochures. A considerable effort is being made by museums to integrate these different means; for example, by creating a coherent graphic layout for digital and printed communication, or by giving the possibility to contact the human helpdesk via either e-mail or chat. The Web site can be designed so as to be reachable or even updateable from visitors inside the museum via touch screen and wireless devices.

But these efforts seem still far from reaching a real, complete integration due to the difficulty of creating a coherent and really usable interface for different means and situations. We have all experienced how difficult it is to integrate different information, coming from different sources, with different formats, inside a common frame like the Web site, and how difficult it is to update it continuously. Moreover, the Web site is simply inaccessible to computer-illiterate persons.

One way to achieve a deeper integration comes from a new generation of natural language recognition systems and their user-friendly interfaces. These applications are able to understand text inputs and spoken language coming from any source (e-mail, chat, Web forms, phone). After getting the input, the system will try to find the appropriate answer by applying complex interpretation rules and by searching different databases inside or outside the museum's infosphere. After the answer is found, it can be transmitted to the user by a variety of means: Web, e-mail, cell phone messages, vocal messages.

Such interfaces can integrate many useful applications: museum mascot, interactive guide, shop assistant, first level help desk, e-learning tutor and customer care. It is also easy to imagine a proactive role; e.g. since the system can dialogue in real time with users, also getting valuable data about their needs and desires, it could personally invite people to museum events, always in a very interactive and natural way.

Is this science fiction, particularly for low-budget museums? Perhaps not. It is now possible to develop powerful and easily maintainable solutions on low-cost platforms, integrating the museum's IT infrastructure with artificial intelligence based characters. They could act as a front end for natural language recognition engines, speech recognition systems, and existing database and content management applications. The paper will present these new solutions, together with the first case histories of natural language interfaces usage. Finally, the paper will explore some of the exciting possibilities opened for museums and cultural institutions by the integration of these innovative means.

Introduction

During 2002 we (Stefania Boiano and Giuliano Gaia) launched a project to rebuild the Science Museum of Milan's Web site (www.museo scienza.org). Our aim was to build a more usable, complete and coherent Web site, and to experiment with new online education tools. The first part to be realized was a very large section about Leonardo da Vinci, with Flash interactive educational games and a massive amount of information about Leonardo and his works (see www.museoscienza.it/leonardo for an Italian preview). A Shockwave 3D version of the Ideal City imagined by Leonardo was also under development by the Politecnico of Milano (Gaia-Barbieri 2001).

During the previous years, some experimental chat sessions between an expert of the Museum and some remote classes had proved successful and involving for pupils (see Gaia 2001). The problem was the difficulty of devoting staff time to this activity. So we included in the rebuilding project the implementation of a chatbot, a piece of software designed for dialoguing on the Internet via text chat with human users. The bot task was to answer visitors' questions about Leonardo and his works.

Chatbots are not a new idea. Eliza, the famous software icon who simulated a psychiatric interview in the Sixties, was somewhat an ante-litteram chatbot.

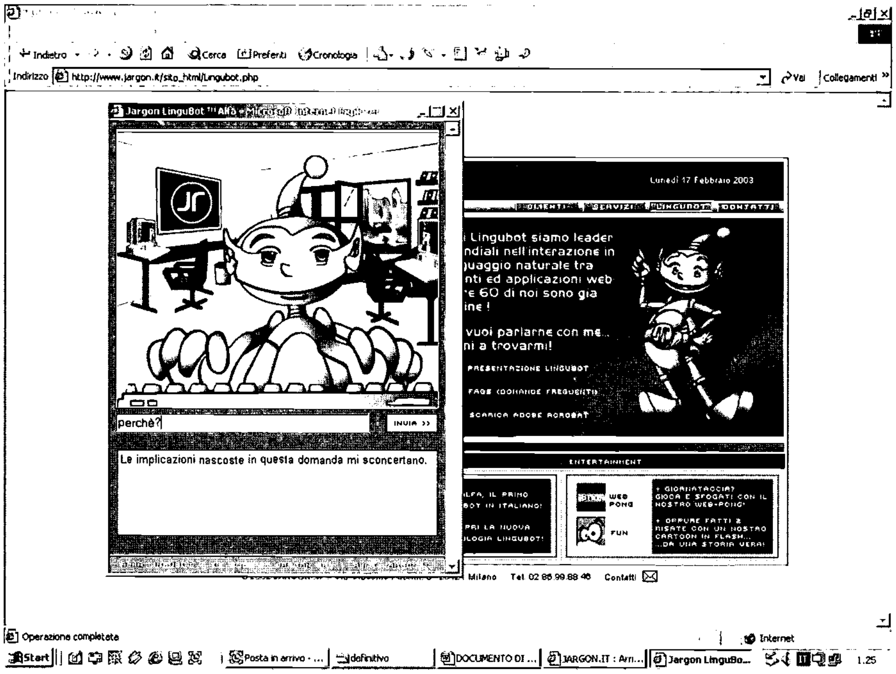

Figure 1: Chatbot created and trained by Jargon – http://www.jargon.it

Chatbot started to flourish in Internet Relay Chat (IRC), an enormous chat environment born in 1988 and still commonly used. IRC offered powerful programming options to its more experienced users, so many programmers started creating software to simulate human chatters (Herz 1995) .Today there are many different kinds of chatbots able to chat in different languages. (Check for examples at www.agentland.com). Most of them are based on the same principle: the text input of the human counterpart is compared to a knowledge base of sentences and keywords, in order to identify a suitable answer using matching rules set by the programmer.

Chatbots are used mainly for entertainment, even if some alternative applications are arising in the commercial field (for example, as virtual shop assistants) and in the cultural one (an interesting example as a virtual tour guide is being presented in MW 2003 too; see Almeida-Yokoi 2003).

In our project, we teamed with two specialists in chatbot technology (Morgana Caldarini, co-author of the present paper, and Andrea Manzoni, both founders of Jargon, a Milanese Web agency). They launched in April 2002 an innovative Italian chatbot answering to the name of Alfa (currently online at www.jargon.it).

The project was quite ambitious and articulated in different steps. Unfortunately, the whole project of rebuilding the Web site was stopped by the Science Museum in December 2002, and with it the chatbot project. So we will describe here all the project steps, even if only the first three steps have been completely or partially implemented.

Step 1: the Bot is Created

The first step is the creation of the bot. This means defining different aspects:.

Application field: as we said before, the first application field we planned for the chatbot was to answer questions about Leonardo da Vinci. It seemed a good start, because:

- It was a way of experimenting with the chatbot extensively in a significant but not risky way (using it as a shop assistant would have been more difficult as a first step).
- An amusing and innovative feature would have been added to the Science Museum Web site
- Since we were creating the large Web site section about Leonardo, it was relatively easy for us to

program the chatbot to push to users specific Web pages related to their questions

While we started with the chatbot only as a sort of virtual expert about Leonardo, we were already planning development of the software towards more innovative applications (e-learning, virtual assistant for online booking and shopping, first-level help desk, and so on). In fact, we saw the chatbot not only as a technological gadget to be added to the Web site, but also as a "natural language interface" in a much wider sense. An interface means that the chatbot can act as a mediator between user inputs and a variety of data sources.

Technology: the idea of using the chatbot as an interface between data sources and the user oriented us towards a technology suitable for integration with multiple data formats and sources. After exhaustive analysis, we decided that the Lingubot technology from Kiwilogic (http://www.kiwilogic.com/) was the most suitable for our project, for the following reasons:

- Great scalability - potentially unlimited number of concurrent user sessions (dialogues)
- Ease of integration with other systems / legacy back-end databases and applications
- Ease of use of the Windows-based Authoring system
- Powerful dedicated scripting language for implementing sophisticated, dynamic pattern recognition rules
- Powerful dialogues logging and statistical analysis system

Target: one of the essential things in every communication project is to define the target carefully. One important feature of bots is that they are able to adapt to the user; i.e. they can change the language and the content of their sentences depending on the user characteristics emerging from the dialogue. Some main coordinates need to be given to the bot; in our case, we decided to focus the virtual guide on youngsters, privileging easy language and concepts.

Language: being a linguistic interface, the choice of the language is very important. Kiwilogic offered Lingubots different linguistic knowledge bases; we chose to start with Italian, developed by Jargon, because this would have allowed us to control idiomatic expressions better and to get the finer details of the user-bot interactions.

Physical aspect: the chatbot can be associated with an image, a flash movie or other file types; for example using 3D plugins or streaming video. This means giving to the bot not only a face, anthropomorphic or not, but also gestures, facial expressions, and so on (coordinated with text outputs). Of course the aspect has to be carefully thought out, because it can deeply influence the user perception of the bot, and therefore the interaction. We decided to avoid the obvious choice of making it look like Leonardo, because it would have been banal and very difficult to realize – how can you in fact program software to talk like a genius?

We decided to make it resemble a Leonardo's machine (a talking flying-screw). This offered us some advantages:

1. It was possible to offer technical information about the way Leonardo's machines worked from an unusual and amusing "subjective" point of view.
2. It was possible to justify some errors or misunderstandings of the bot during the dialogue due to its artificial, not human, nature :"Leonardo did not teach me about that, so I cannot understand what you are trying to say me..."
3. We had more freedom to create new "gestures" – a machine can nod and smile, but also fly, transform itself, and so on, in a more natural way than an anthropomorphic image.

Personality: strictly connected to the physical aspect is the character personality. To make a character realistic you have to create not an "answering machine" but something simulating the complexity of a human relationship with attitudes varying according to the conversation. For example, if a bot is insulted by a human being, to be realistic it has to become angry and stay so for a significant amount of time. Facial expressions and gestures can reinforce its personality. So it is important to define the personality well and to shape answers accordingly, even if a certain degree of incoherence is also necessary not to make the bot reactions too predictable. We tailored for our bot a calm and solid behaviour with some irony.

Step 2: the Bot Learns

Once the bot is defined, the difficulty begins: in order to talk, the bot has to learn how to interpret and answer the user's input. As said previously, a Lingubot comes with a pre-set knowledge base of the specific language and a corpus of interpretation rules, but this is only the first floor of a much more complex building. The bot has to be carefully "trained" to give it the possibility of accomplishing its designated task. In our case, as a virtual guide, the bot had to learn a lot about Leonardo and his life. Moreover, the most difficult task is to connect questions to a specific answer and to understand when the user is going outside the bot field of expertise.

We had the valuable help of two assistants, Davide Radice and Felice Colasuonno, both coming with a strong humanistic background but not having special programming skills nor previous experience in chatbot programming.

The bot was trained by going through 4 different phases, which overall lasted one month:

1. The two operators were trained by experts at Jargon, and learned all the techniques and concepts behind making a Lingubot. This required one week of lessons and hands-on experimenting.

2. After the training session, the operators analysed a huge corpus of information about Leonardo coming from our Web site, defined a grid of concepts and their relations, and then re-wrote the content of the corpus in dialogic form.

3. The specific knowledge base was loaded into the bot's "mind", by creating all the required special dictionaries and the generalised pattern recognition rules (called "recognitions"), and by linking the proper answers and bot "emotions" to each recognition. Then the specific knowledge base was integrated and refined using interviews with museum guides, FAQ's lists, e-mail from visitors of the museum Web site and log file analysis. During this step, the Lingubot was tested by our actually talking to it and checking that the recognition of concepts worked independently from the words used to express them.

4. The Lingubot was activated and published in a restricted area, to be tested by a larger number of beta-testers: meanwhile, the Lingubot was integrated with the Web site in order to open pages, and show documents and animations about Leonardo and his inventions.

Step 3: the Bot Chats

After receiving basic training, the bot is able to start written conversations with real users in the predefined language. To understand how this happens, let's see how the bot reacts to a user question

1. As a first try, the bot compares the sentence to its knowledge base, searching for a match and the related answer. On finding it, the bot writes back the answer; otherwise, it passes on to the next step.

2. If no perfect match is possible, the bot tries to guess the meaning of the sentence by identifying some significant keywords, matching them with grammatical rules and the general context of the conversation (for example, if the previous questions of the user were about a specific Leonardo's manuscript, in the question "where can I find it today?" the bot assumes that "it" is related to the manuscript). This assumption ends when it clearly realizes that the user has changed subject.

3. If the bot definitely does not understand the meaning of the question, it can:

   a) admit that the input has not been understood and ask the human counterpart to reformulate the question

   b) try to deceive the user with conversational tricks, such as changing subject, or telling a joke, or "nodding"... as every one of us has done in similar situations!

Of course, the probability of getting the right answer increases with the size of the data base. Hence, continuous work on widening and fine-tuning the knowledge base is necessary: a great help with this regard comes from the analysis of the users' dialogues.

The interpretation of the meaning of human sentences is a most problematic point. For this reason Turing chose the man-computer dialogue as its operative test for Artificial Intelligence – even a last-generation bot would not pass it, since chatbots can always be put on the wrong track by ambiguities of the language in the sentences of the human user, or just by typos. Nevertheless, a well-trained bot can carry on even long conversations in a credible way and efficiently perform many tasks.

Once properly trained, the bot is ready to start its life as a virtual expert, since it is able to "understand" most user inputs and try to find appropriate answers. The bot can also open pre-defined Web pages related to a question, and it can recognize returning users (using cookies) - offering the possibility of a high degree of personalization. When a user is recognized, the bot can scan logs of previous dialogues and tell the user: "I remember you were interested in Leonardo's flying machines. Do you know we are opening a special exhibition on this subject?"

Logs generated by dialogues are valuable because they:

- Provide much deeper insight into users' needs and desires than do usual Web site statistics
- Permit one-to-one marketing activities like the one described before
- Offer the possibility of refining the bot's abilities to understand human sentences and to fix weak points or answers that prove to be badly accepted by users (for example, an ironic answer felt as insulting by most interlocutors)

Like humans, bots never finish learning. By studying conversations made by the bot during its activity, operators can fix conversational situations in which the bot does not give a proper answer, and can identify subjects human users are particularly interested in and enlarge the knowledge base accordingly. In fact, most of the knowledge base work is made after, and not before, the online publishing of the bot. At least six months of continuous refining work is usually needed to make the Lingubot able to carry on a good number of conversations on its specific subject.

For instance, Alfa, the Jargon bot, in the period April-December 2002 had more than 200,000 conversations. All of these conversations were recorded and analysed, both manually and automatically, by Jargon Authors. This enabled the inclusion in the Italian Lingubot Basic Knowledge Base of more than 40,000 new terms, 4,500 recognition patterns and thousands of keywords and interpretation rules. After 9 months of extensive training, Alfa's recognition success rate is now stable at around 97%.

Unfortunately, as mentioned before, in December 2002 the Leonardo project was stopped, so the bot could not undergo the extensive testing scheduled before the launch, set for February 2002. From here on, we will describe the further steps of the project which were planned for a period of 9 months from the launch.

Step 4: the Bot Sends e-mail

Once the bot recognizes a user, it can also offer the possibility receiving e-mails about specific subjects. These emails can be just traditional text newsletters, with an extra link to the bot. By clicking that link, the user can start a conversation with the bot directly talking about the newsletter content. For example, the user can ask the bot why it signalled that specific conference to him, or ask questions on specific things that are not specified in the conference Web page, or ask about other events related to that one... In fact, a well-trained bot added to a newsletter can enhance significantly both user satisfaction and the collecting of feedback.

During our work on the Leonardo section, we created a virtual postcard page where the user could send to friends unusual drawings by Leonardo, together with a message. We experienced a high level of interest and satisfaction with this page. We planned to offer users the possibility adding to their postcards a link to the bot. This way, the users could offer friends the possibility starting an unusual conversation about the postcard image. Moreover, the users could tell the bot secret messages to be revealed during the conversation with the friend. This was intended as a tool to increase Website visits, but it also could have educational side-effects: for example, the bot could focus conversations on the image, stimulating the user with cultural questions; only after getting a right answer would the bot reveal the message of the user's friend.

Step 5: the Bot Talks

Downloading a free plugin, Web site visitors are able not only to read what the bot says but also to hear its words, thanks to a voice synthesizer. We decided to implement this feature because it made conversations more engaging, especially for pupils, and also more accessible to visually impaired people not equipped with vocal browsers. The main problem with vocal synthesizers is intonation: since synthesizers cannot understand the meaning of the sentence, they pronounce every word singularly, with no connection to the other words. This is a major problem especially with languages like Italian which are strongly based on intonation (for example, interrogative meaning to a sentence in Italian is mainly given by intonation). The choice to make the bot look like a machine was also intended to "justify" its unnatural way of talking.

During 2002, AT&T released a new generation of voice synthesizers called "Natural Voices" http://www.naturalvoices.att.com/ which are able to read English with nearly natural intonation, and able to give "colour" to sentences. This could improve significantly a bot's talking capabilities.

Step 6: the Bot Listens

Talking is only half of the vocal interaction, and the easy one. What we really wanted was to make our bot able to listen to the users on the Web and outside the Web. In our opinion, this step is very important because bots could be very effective if used inside the museum, for example in a multimedia kiosk.

During 2002, a renowned Italian group of digital artists, Studio Azzurro, tested in the Science Museum an innovative system called "Torkio". Torkio is a digital character interacting with museum visitors, remotely controlled by a human operator. The success of the installation is high, but the need of skilled operators limits its operating times. Observing the high emotional impact of the digital character on museum visitors, we thought it would have been very interesting to put the bot in the museum too; the bot could have been an effective virtual guide. A "talking kiosk" could invite visitors to play, could ask questions, offer answers, show multimedia files related to their questions, collect their emails for newsletters, and be an amusing technological exhibit itself.

In this case the limit is technology. Till today, speech recognition systems work well only when:

1. the system has been trained a lot on a single voice (like IBM Via voice system, that requires long training by the user before becoming really usable)

2. the system has to recognize only few words (like recognition systems of phone companies' directories: they have to recognize only single words clearly pronounced by the user).

None of these is true in our case, since the bot should be able to interact with plenty of new users who pronounce long sentences in a natural way, sometimes "worsened" by the emotions (e.g. a laugh or anger) created by the conversation, and with a lot of environmental noise. But a new generation of low-cost and effective speech-recognition software has been announced by DARPA-funded researchers at Carnegie Mellon University http://www.speech.cs.cmu.edu/. The English and French speaker independent speech recognition modules are already available, and the Italian module is being developed by Jargon. So it is possible that an effective listening ability could soon be added to chatbots.

Step 7: the Bot Goes Wireless

Going out from the Web does not mean only going into the museum as a new-generation interactive kiosk. Kiwilogic provides Lingubot with the ability of interacting with users using cell phones messages (SMS), and more generally using wireless systems.

This opens up interesting application fields:

1. answering users' questions about the museum, its activities and its collections both inside and outside the museum. A user could send an SMS to the bot to know when the conference of the day is starting, or if a certain book is available in the museum library.

2. Many museums are experimenting with handheld devices, as witnessed by the rich "handheld" session in MW2003 (see, for example, Cigliano-Monaci 2003). A bot could easily interact with the user through a handheld device, asking the user to find certain objects inside the museum, in a sort of interactive treasure hunt.

3. If equipped with an effective voice recognition system, the bot could talk on the phone with users, providing much broader access to its information and interactions, even to computer-illiterate people.

Step 7: the Bot Is Integrated with the Museum Data Sources

Lingubot technology can be easily integrated with external databases (via ODBC/JDBC) and to external applications and sources using HTTP or TCP/IP. Bots can also call applications directly. This means that the bot can get data from any digital source inside or outside the Museum infosphere.

In our project, we had planned to integrate the chatbot with the following data sources:

1. collections database

2. library database

3. guided visits booking system

4. Web site

5. internal bookshop

6. online Web international bookshops

This way the bot would have been able to provide different services to the user:

- provide information about collections, links and books related to them

- act as a "virtual librarian"

- be a promoter and assistant in booking guided visits

- be a promoter and assistant in online and offline shopping

The key and the strength of the system was the integration of the different data sources. The bot could in real time adapt its messages to the situation of the bookings; for example not promoting events already full or books sold out. Of course, the bot is only a first-level help-desk; if the bot is not able to fulfil its mission, for example because it does not understand what the user is asking, then it redirects the user to the human help-desk.

Step 8: The Bot Extends Itself by Interacting with Its Peers over the Web

Integration could go further than the previous point. Since a well-trained bot becomes a specialized database itself, and it is easy to make bots communicate between themselves over the Internet, it is also easy to imagine a situation in which a bot in the Science Museum of Milan gets a question on English steam trains and connects to its "colleague" at the British National Railway Museum to find the answer... this would be even easier for more organized tasks such as online bookings or shopping, and could allow us to develop one of the first effective, practical working infrastructures of the "semantic Web".

Conclusions

Every time a new technology hits the market, it has to answer three main questions posed by potential users (including museums):

1. **Is it useful?** In the paper we have described the plurality of applications and services that this technology could help to create inside and outside the museum. In our opinion the key point is to

see this technology as a new interface, capable of integrating and enhancing traditional Web and digital interfaces.

2. **Is it expensive (in terms of money and resources)?** Since many of the necessary software packages are open source, the overall budget for hardware and software is not high. The system requires a significant amount of work for training and refining, but this part can be a very interesting task, for example, for university students and researchers. We discovered that training a bot can be a very stimulating and involving task even for teenager students. This adds further educational value to the tool.

3. **Is it lasting?** In a fast changing environment like the Web, with new technologies rising every day, it is difficult to predict whether a technology will last or will be ephemeral. Chatbots, however, go in the direction of simplifying the interface and making it more natural for the user, and in our opinion this should make their technology a good bet, or at least an experiment worth trying.

References

Almeida P., S. Yokoi (2003). Interactive Conversational Character as a Virtual Tour Guide to an Online Museum Exhibition, paper presented at Museums and the Web 2003. http://www.archimuse.com/mw2003/abstracts/prg_200000698.html

Cigliano E., S. Monaci (2003). Multimuseum: a multichannel communication project for the National Museum of Cinema of Turin, paper presented at Museums and the Web 2003. http://www.archimuse.com/mw2003/abstracts/prg_200000703.html

Herz, G.C. (1995). Surfing on the Internet, Little Brown & Company.

Gaia, G. (2001). Towards a Virtual Community, paper presented at Museums and the Web 2001, http://www.archimuse.com/mw2001/papers/gaia/gaia.html

Gaia G., T. Barbieri (2001). HOC - Politecnico and Museum of Science and Technology of Milan: A Collaborative Exploration of Leonardo's Virtual City, paper presented at ichim 01, http://www.archimuse.com/ichim2001/

Interactive Character as a Virtual Tour Guide to an Online Museum Exhibition

Pilar de Almeida and Shigeki Yokoi, Nagoya University, Japan

http://www.mdg.human.nagoya-u.ac.jp

Abstract

In real museums, visitors may seek the help of a human tour guide in order to guide them through the exhibition and present information tailored to their interests. Accordingly, online museums could also benefit from digital "life-like" characters in order to guide users to virtual tours and to customize the tour information to users' interests. Digital characters have been explored in online museum Web sites with different degrees of interaction and modes of communication (Bertoletti et al., 2001; Adams et al., 2001; Doyle, P. and Isbister, K. 1999). Such researches, however, do not explore interactive "conversational" characters. Our main interest is in a tour-guide character that provides tour information through dialogue with the users. Our research developed an interactive conversational character that establishes a "character-user" dialogue while guiding the users through the virtual tour. The tour-guide character provides tour information and responses to users' textual inputs by speaking through audio output and gesturing accordingly. Moreover, the tour-guide character attempts to provide adaptive guidance, to perform engaging storytelling and to promote users' participation through the virtual tour. In this paper, we describe the experiment focusing on the systems implementation, the dialogue creation process and the character's ability to detect users' interests. We further describe and analyze the system's evaluation. We hope lessons learned could serve as a base for developers of interactive characters as tour guides.

Keywords: virtual guide; adaptive guidance; virtual tour; interactive virtual character; storytelling

1. Introduction

In real museums, visitors may seek the help of a human tour guide to guide them through the exhibition and present information tailored to their interests. Accordingly, online museums could also benefit from digital interactive "life-like" characters in order to guide users to virtual exhibitions and to customize the virtual tour information to users interests.

Digital characters have been explored in online museum Web sites with different degrees of interaction and modes of communication (Bertoletti et al., 2001; Adams et al., 2001) However, these researches have often limited their use to embodied alternatives to traditional menu or prompt-driven mechanisms for performing simple tasks such as searching for files or providing help information embedded in Web pages. (Bertoletti et al, 2001; Adams et al 2001). Closer to our goals is Isbister's research where an agent guide tracks users' words in a chat environment and provides more detailed or less detailed tour information accordingly (Doyle, P. and Isbister, K. 1999). Our interest is, however, in an interactive conversational character that establishes a customized person-to-person dialogue while guiding the user through the virtual exhibition.

In recent years, due to technological advancements, interest in interactive conversational characters has sharply increased mainly in response to e-market strategies. E-service providers explore their potential as sale assistants that perform intelligent-like behaviors such as greeting users, answering questions and suggesting personalized products. Some famous examples are Linda (by Extempo), Julia (by Virtual Personalities), Nicole (by Native Mind), Lucy Mcbot (by Artificial Life). Such characters not only instruct, guide, advise and perform tasks on users' behalf, but also enhance the communication with users through both verbal and non-verbal interactions such as hand gestures, facial expressions and gaze movements.

Likewise, research on conversational characters in learning environments has concluded they not only facilitate the user experience with the learning environment but also greatly improve the computer's ability to engage and motivate students (Lester, J. et al 1997).

Traditional complaints about interaction with conversational characters highlight that their reactive, context-free conversation and their lack of goals and motivation for interaction inevitably lead users to interact for a short period of time as well as increase the potential for unmet expectations regarding the character's intelligence. Thus, researchers on character development have focused attention on the design of motivational structures, emotional and personality traits and behavior control systems for characters to perform in context-specific environments, with well-defined goals and social tasks (Hayes-Roth 1998; Doyle, P. and Isbister, K. 1999; Lester, J. and Rickel, J. 1999). Emphasizing the importance of the social context for shaping user-character interactions, Doyle has suggested that in social tasks, in contrast to mechanical tasks, uncertain goals or goals of satisfaction rather than correctness, and dialogue rather than command-structured interactions, shape a good domain where an interactive conversational character may work well (Doyle, P. 1999).

Within this framework, our research attempts to shape dialogue interactions between an interactive conversational character and the user in the specific context of a guided tour to a virtual exhibition. We hope to examine features of conversation tracking and interest monitoring in order to improve the dialogue process and to identify some features that could serve as a framework for developers of interactive characters as tour guides.

By observing the qualities of a human tour guide in a real museum, we developed a virtual tour guide to a virtual exhibition. A human tour-guide is expected to provide general information about the tour, to highlight curiosities about its objects, to tell stories related to the tour and to answer most frequent questions regarding tour and exhibition topics. Moreover, a tour guide's performance must include abilities to provide adaptive guidance and engaging storytelling, to stimulate visitors' participation and interest in the exhibition and to reincorporate information from previous tours.

Considering such qualities, we developed an experiment with an interactive character as a tour guide to a virtual exhibition. In the experiment, the user explores a virtual XVI Century Portuguese ship and objects inside while talking with and receiving information from the tour guide. In this paper, we describe the experiment focusing on the system implementation, the dialogue creation process and the character's ability to detect users' interests. We further describe and analyze the system's evaluation.

2. Guiding an Exhibition Tour

In real museum exhibitions, curators display objects, technical information, and narrative in space in order to convey the exhibition message (Roberts, L. 1997). Tour guides often reinforce such narrative by conveying information that can be roughly divided in two levels: a presentation level, conveying general contextual information about the group of objects and their meaning, and an object-specific level, conveying information about the displayed object such as how it was made, by whom, why, when, how it was used, etc.

Presentation level information is normally introduced prior to the viewing of objects, at the beginning of the tour and/or at the entrance of each exhibition room to provide context. Human-guides often present such information in a compulsory, non-interactive way in order to establish the atmosphere for conveying the exhibition message.

For object-specific information, on the other hand, human-guides are more susceptible to visitors' interests. They are more likely to adopt a reactive approach in which the content of their speech will depend on how visitors interact with the exhibition material.

Moreover, besides these levels of information flow, a human tour guide must also have engaging storytelling abilities to stimulate visitors' participation and exploration of the exhibition environment. Telling about curiosities and calling visitors' attention to interesting details are some of the techniques used to shape an engaging guidance.

3. The Experiment

Our experiment developed a virtual tour guide that leads a tour to a XVI Century Portuguese ship. Users are invited to visit four rooms of the ship in a predefined sequence (Figure 1). Inside each room, room-related objects are exhibited. (Table 1) The rooms are HTML pages with images of objects (Figure 2).

Room	
1. The Deck	Objects: a sailor, a cabin-boy and a toilet basket
2. The Pilot Room	a compass, an astrolabe and the Pilot
3. The Storage Room	barrels of wine, biscuits, water and ammunition
4. Captain's Room	the Captain, a route map and the Cross of Christ

Table 1: Rooms and exhibited objects

The virtual tour guide provides information about each room and its objects. As the user enters the room, the virtual tour guide presents the room telling how it was used and what kind of objects the user can find there. Next, the user explores the objects in the room by looking at their images and reading their detailed textual information in a small window that pops-up when the mouse is over the object. After that, the tour guide makes some comments about what the user saw. Such comments will serve to start a dialogue that gives the tour guide tips on the depth of the users' interest in the room's objects and topics.

Topics for the guide's comments and stories were generally divided according to the room and the objects exhibited inside. In this way, the "Deck" would lead to conversation and stories related to sailors' sanitary conditions, crew and food. The "Pilot room" would lead to conversation and stories related to sailing instruments, the helm, pilot responsibilities, etc. The "Storage room" covered information on food, armaments, diseases, animals on board, etc. And, finally, the "Captain's room" would cover information on the captain's responsibilities, the objectives of the discovery voyages, political and economical interests of the time, Portugal and India, etc.

3.1 System Description

The virtual tour guide system consists of two components: a computer keyword spot program with a natural language processing system that delivers pre-scripted statements from a knowledge-content database, and an MS-Agent character that performs such statements as gesture-choreographed dialogue pieces.

The tour guide's remarks are delivered both, through textual balloons and audio output. Users' inputs were textual entries in a textbox. (Figure 2)

The dialogue plan and the knowledge-database were designed and scripted with ImpCharacter Editor 2.1 developed by Extempo Systems Inc. (http://www.extempo.com) (Figure 3).

3.2 The Tour Guide Knowledge-Database

The tour guide knowledge-content database stores the dialogue steps, keywords (of three types: general keywords, room specific keywords and after-dialogue keywords), pre-conditions for the delivery of dialogue pieces, as well as the tour-guide's personal history and mood settings.

By using pre-condition features, the tour guide can personalize the dialogue by tracking each dialogue piece and visited room in order not to repeat given information. In this way, even if the user asked to go back to a specific room, the tour guide would not offer the same dialogue pieces. Pre-conditions were also valuable for triggering extra stories whenever the user was interested in more detailed information.

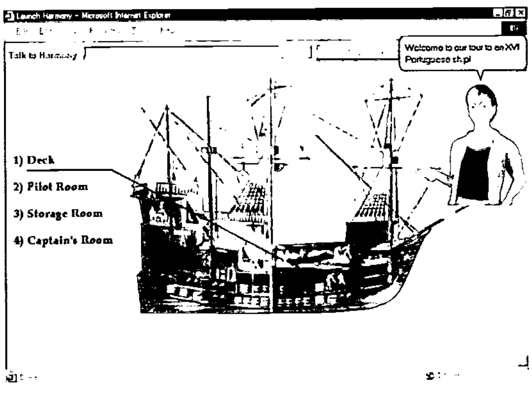

Figure 1: Beginning the Tour

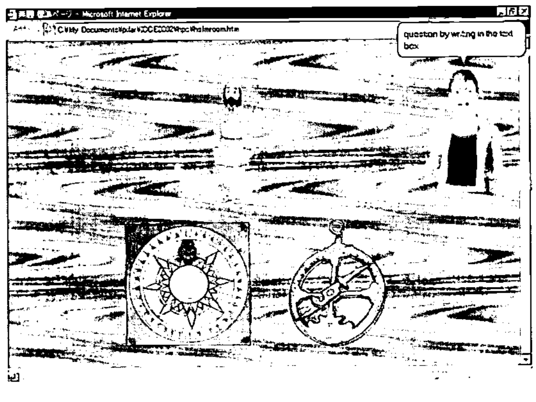

Figure 2: Screen of the Pilot room

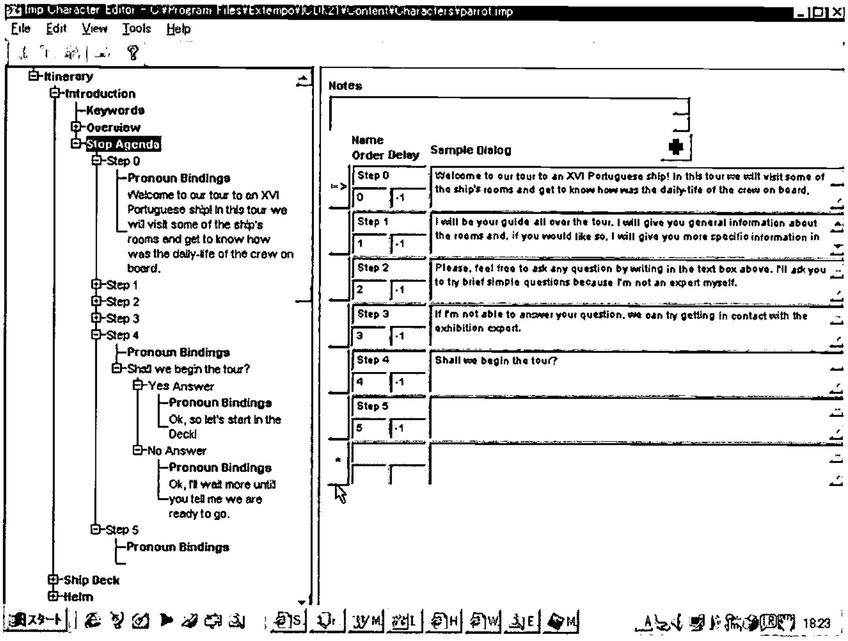

Figure 3: Authoring the Guide's Dialogue and Knowledge

Also, the ability to arrange keyword schemes allowed the tour guide to make some "intelligent" inferences. For example, through the schema: [food] = [biscuits, bread, lunch, dinner, meat, fruits, eat], the following dialogue is possible:

> *Visitor: What would sailors eat for dinner?*
>
> *Guide: We will get detailed information on the sailor's food as soon as we get to the storage room!!*

3.3 Dialogue Creation and Detecting Users' interest

Planning the dialogue flow for the tour-guide was the main work of our experiment. We wanted the guide to give information following the exact flow of a tour and exhibition visit. Moreover, we wanted the guide to detect users' interest in an indirect way. We feared that generic questions such as "Are you interested in this?" Or "Would you like more information about that?" would not sound natural and would not differ from a traditional multimedia menu-like environment.

The key moment to detect users' interest was during the dialogue triggered by the guide's comments after the user had explored the room. As we said, according to user's reactions to the tour guide comments, the tour guide would provide (or not) extra-stories about the room. The comments are, thus, supposed to be very provocative. Some examples of comments are:

> *"Did you see the rope just beside the toilet basket? Do you know what it was used for?" (Deck)*
>
> *"Do you like olive oil and vinegar?" (Storage Room)*
>
> *"Can you guess any spice that could be found in Indian markets?" (Captain's Room)*

During the dialogue triggered by the tour guide comments, the tour guide checks for some supportive words (such as "nice", "cool", "interesting", "really?", "you bet!", "definitely", "sure",...) according to the dialogue piece. On finding such words, the tour guide assumes the user is interested and provides more information on the room.

The piece of dialogue below illustrates how indirect the detection of users' interests is.

> *Guide: Did you see the rope just beside the toilet basket?*
>
> *Visitor: Yes*
>
> *Guide: Can you guess what it was used for?*

Visitor: It seems a brush

Guide: Yeah. This rope with a frayed end was used as toilet paper. It was kept under water and pulled up when someone needed it!

Visitor: Yuck!!! I would never use it!!

Guide: Yeah, people were not very clean those days

In this example, the supportive word is "yuck" (along with several possible others such as disgusting, terrible, dirty, God, etc). The guide assumes the user had a good reaction to this story and will provide further stories.

Another way to detect users' interest is from users' questions. Keywords for possible questions and their answers are pre-defined. Each question is related to a specific room. Whenever the user asks a question, the guide provides the answer and tells extra stories when the user reaches the room related to the question.

4. Evaluation & Analysis

We conducted preliminary tests to our experiment with eleven university students. We used both the tour's dialogue transcripts and user evaluation questionnaire data to determine how effective the tour guide was as a way to convey adaptive guidance and storytelling to virtual exhibitions.

In the evaluation questionnaire, students were asked to grade the experience with the tour guide using a scale from "very good" = 4 to "not good" = 1.

As shown in Table 2, in regards to an entertaining experience, the average value suggests that visitors found the experience of dialoguing with a virtual tour guide character really enjoyable (Avg=3.38). Data shows also that the tour guide was effective in motivating users to explore and learn more about exhibition topics (Avg=3.3). Yet visitors were slightly less sure about the extent to which the character's remarks helped them better understand the exhibition topics (Avg=3.15).

In regards to the quality of dialogue interactions and the tour guide's ability for adaptive guidance,

Questionnaire Evaluation Results	**Average Score** 1 = not good / 4 = very good
Entertaining Experience	3.381
Learning Experience	
Motivation	3.3
Better Understanding	3.15
Dialogue Interactions & Adaptive Guidance	
Encouragement to speak	3.05
Adequate responses	2.35
Interests detection	2.15

Table 2: Evaluation Results

users felt the character satisfactorily encouraged them to speak and interact with exhibition material (Avg=3.05). Yet, regarding the character's ability to respond to users' questions, the average achieved a regular level (2.35). This data suggests that more work should be done to improve the character's knowledge; that is, more keywords to its database. Users also hardly recognized the ability of the character in providing further information when they were interested (2.15). This data shows that our decisions for capturing users' interests (based on "supportive words" and questions) were not perceived as the character's natural ability to provide further information. This could be highly alleviated by language tips such as "Since you asked about….", "Since you seem interested in…". In our experiment, such introductory remarks were nonexistent. The tour guide simply provided extra stories when she detected users' interest. Also, users easily confused this programmed choice with the tour guide's inability to respond to certain questions.

The last question of the questionnaire was an open question regarding the presence of the character. The question was: "If you could choose the same virtual exhibition with or without the tour guide, which one would you choose?". All users with no exception chose *with* the tour guide. In general, users reaction to the experiment showed the experience of talking to the tour guide was highly different from individual multimedia exploration, and it played an important role in motivating them to explore the exhibition material.

5. Lessons Learned & Future Work

As a result of the evaluation process, some preliminary (and relatively simple!) corrections were evident and would highly increase the tour guide's performance. We hope they may assist future developers of tour guide characters.

- **More knowledge in her knowledge–database, especially exhibition-related vocabulary.** The experiment shows that the universe of questions is highly limited to the exhibition information. Though hard work, it is easier than one might think!

- **Better use of personal introductions.** Her personal introduction and personal questions to users could be used to "break the ice" and to determine users' interests.

- **Tension alleviation when she does not understand questions.** She must provide a way out of the tense cycle created whenever she does not understand questions. Rather than only ask to rephrase or say she is sorry, she must offer a different conversational topic.

- **Users' prompt understanding of an "exhibition-like" information flow.** Users should be told they are about to experience the context of an exhibition visit and that they will receive information and are expected to behave accordingly.

- **Clear names or numbers for exhibition objects.** Users may want to ask about objects and they need clear ways to identify them.

We are now making the corrections listed above and we plan to conduct further tests that may compare students' interactions with the virtual exhibition in environments with and without the virtual tour guide.

Conclusions

In this research we developed an interactive character as a virtual guide to an online museum exhibition. Features such as adaptive guidance, engaging storytelling, stimulation of users' participation were attempted while building the character's knowledge and dialogue creation process. Though evaluation results show some limitations of the tour guide as a successful dialogue-builder, we believe this was mainly due to the experiment's preliminary stage and could be substantially alleviated with the corrections listed in section 5. Moreover, despite such limitations, the experiment successfully proved the potential of conversational interfaces to motivate users to interact with multimedia material. We hope lessons learned may serve as a foundation for developers of interactive characters as virtual tour guides.

Acknowledgements

We acknowledge the support of the "Foundation for Fusion of Science & Technology".

References

Adams, C. et al (2001) Bringing Curatorial Process to the Web. In *Museums and the Web, 2001* USA, Seattle.

Almeida, Pilar (2000) "A Prototype System Development of an Interactive Narrative Model for Online Museums.Åh Journal of the Japan Information-Culture Society *Vol. 7 No. 1, pp.73-80.*

Bertoletti, A.; Moraes, C.; Costa, A. (2001) Providing Personal Assistance in the SAGRES Virtual Museum. In *Museums and the Web, 2001* USA, Seattle, 2001.

Booth, Ben (1998) Understanding the Information Needs of Visitor to Museums. *In Museum Management and Curatorship*, Vol. 17, No. 2, pp. 139-157.

Doyle, P., and Isbister, K. (1999) Touring machines: Guide agents for sharing stories about digital places. *AAAI Fall Symposium on Narrative Intelligence.*

Doyle, P. When is Communicative Agent a Good Idea? In *Workshop on Communicative Agents of the Third International Conference on Autonomous Agents*, Seattle WA, May 1999.

Hayes-Roth, B. (2001) Adaptive Learning Guides *In Proceedings of IASTED Conference on Computers and Advanced Technology in Education*, Banff, Canada, June.

Hayes-Roth, B.; Doyle, P. (1998) Animate Characters. *In Autonomous Agents and Multi-Agent Systems*. Netherlands: Kluwer Academic Publishers. I, 195-230, 1998.

Lester, J.; Converse, S.; Kahler, S.; Barlow, T.; Stone, B.; and Bhogal, R. (1997) The Persona Effect: Affective Impact of Animated Pedagogical Agents. *In Proceedings of CHI'97 (Atlanta GA, March*, ACM Press.

Lester, J.; Rickel, J. (1999) Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments. *In International Journal of Artificial Intelligence in Education, 2000.*

Microsoft/ MSAgent Library (www.msdn.microsoft.com/library/en-us/msagent/).

Roberts, L. (1997) *From Knowledge to Narrative: Educators and the Changing Museum.* Washington: Smithsonian Institution Press.

Consulted Web sites of Portuguese History:

"Na Crista da Onda" (http://www.cncdp.pt/cncdp/crista/)

"Projeto Memoria" (http://www.projetomemoria.art.br/)

Experiencing Art on the Web with Virtual Companions

Ido A. Iurgel, ZGDV e.V. (Computer Graphics Center), Germany

http://www.zgdv.de/zgdv/departments/z5/index_html_en

Abstract

Experiencing artworks in the Web aggravates the problem of the absence of any historical, cultural and social context, because the Web is a veritable nowhere. This paper stresses the importance of social and narrative access to art, and presents an interactive group of virtual characters as an alleviation of this problem. The benefits and overall concepts of this approach are examined. Particularly, the importance of guiding the virtual characters to establish close emotional relations with the user for achieving narrative immersion is stressed, and ways of accomplishing this task are sketched.

Keywords: Virtual Characters, Hybrid Discussion Group, Interactive Storytelling, Emotional Computing, Social Computing

Introduction

Viewing artworks on the Web is an experience absent of any historical, cultural, or social context. While this problem is well known to traditional museums, the Web, as a veritable nowhere, aggravates the situation. Consider someone looking at a Byzantine icon at a Web museum with a home monitor. A visitor – henceforth the "user" - not very well acquainted with the subject is likely to miss much of the meaning of the artwork, and the art experience will be shallow and uninteresting.

It is certainly helpful to offer some additional written explanation or even video material, but still, the experience remains intellectual and solitary, rather than emotional and culturally enriching. Additionally, it is not easy to motivate the visitor to study the additional material, instead of clicking away after some superficial glance.

What can be done to provide a more involving and enriching art experience to the user? According to the European hermeneutical tradition, understanding of an art work is inherently a social, dialogic and emotional process (Cf. Gadamer, 1993; Heidegger, 1973). There is no absolute meaning, but only restricted perspectives that belong to a whole live world. In order to understand a Byzantine icon, a user must understand the artists who made it and the worship of the believers, their way of living, their feelings, and their convictions. This is not simply a matter of increasing a person's knowledge, but of fusing his own horizon with the horizons of the people to whom the icon had an essential meaning. This is not only an intellectual but also an emotional process of feeling empathy and identifying with particular individual motivations and biographic experiences.

At this point, narration enters the scene as a means of promoting identification, empathy and understanding. Through stories, it is possible to boost the fusion of horizons in art experience. Examples of this process in traditional media are abundant, e.g. historical novels or films.

A new approach is "interactive storytelling". Interactive storytelling is the discipline that deals with the run-time creation and adaptation of stories, observing emergent constraints like global duration of the interaction or user choices and interests. It aims at combining the narrative immersion of traditional story telling with the advantages of user behaviour adaptation, employing, for example, educational strategies.

In this paper, I will present a narration driven, interactive group of virtual characters with whom the user can fraternize and discuss each other's art experience and perspectives. This group allows reconstituting the contingent meanings of art and regaining the missing emotional and existential perspectives on it by preserving the narrative horizons. At the same time, the group involves the user through interaction and adaptation.

Figure 1: Two virtual characters speaking about a religious icon

This interactive group is being developed in the projectArt-E-Fact (Cf. www.art-e-fact.org). This EU-funded project, a collaboration of several institutions of four European countries, aims at developing generic tools for artists to create dialogic, narrative worlds inhabited by personality-rich virtual characters. The interaction with these worlds is enhanced by specialized utilities like exploration tools that allow the user to uncover different painting layers. Though the project does not aim specifically at Web based use, incorporating an internet version poses no technical problem. AVRML viewer, together with some small plug-ins and a freely available speech synthesizer, are the only technical requirements. The initial demonstrator deals exactly with a museum presentation of the Byzantine icons mentioned above (Cf. www.monastery-artdiagnosis.gr for a link to one partner).

The next section gives an overview of previous work. Then, I will present the overall architecture of the system. After that, I will explain the benefits and some ideas behind the project. I will stress especially the importance of automating the process of creating social and emotional bonds with the virtual characters to foster the experience of narrative immersion.

Related work

The importance of narration for the structuring of experiences is also stressed in psychology, e.g. in the work of Shank andAbelson (1995) and Graesser and Ottati (1995).

André et al. (2000) developed a non-interactive presentation group without narrative structures. Their emphasis lies in natural language generation independent of personality traits of the virtual participants. Mateas and Stern (2000) are developing a very promising system for interactive storytelling. In contrast to our approach, they do not aim at education and presentation, which is probably the reason for the different architectonic and conceptual details, since there is no need to present information in their system, and therefore a less centrally guided approach is possible.

Rickel et al. (2002) are engaged in an ambitious project that comprises natural language interaction, emotional and social simulation, but not information presentation and narration. A. Paiva, I. Machado and others have developed various interactive applications with storytelling and/or tutorial components (cf. Machado et al., 2001, Paiva, 2001), though none of these systems uses interactive discussion for multifaceted edutainment.

Other projects in which our own group is involved are "Geist" (cf. www.tourgeist.de), where interactive storytelling is used, but without any natural language interaction, and "Virtual Human", which aims at the creation of a narrative tutoring system with talking virtual characters (cf. www.virtual-human.org).

Architecture of the System

The basic architecture is depicted in Figure 2. The information process is centralised and hierarchic, the flow of control starting at a story engine and passing through a scene engine to the dialogue manager or directly to the scene player. The player executes the scripts delivered either by the scene engine or by the dialogue manager and controls the partially autonomous virtual characters and objects. The current architecture separates strictly interactive and non-interactive episodes. In an interactive episode, the user can talk to the virtual characters using a keyboard, and the dialogue manager is responsible for the generation and distribution of adequate answers. In the case of a non-interactive episode, the scene engine controls the virtual characters directly. Non-interactive episodes are fixed sequences authored in advance, much like the dialogues in a traditional theatre play.

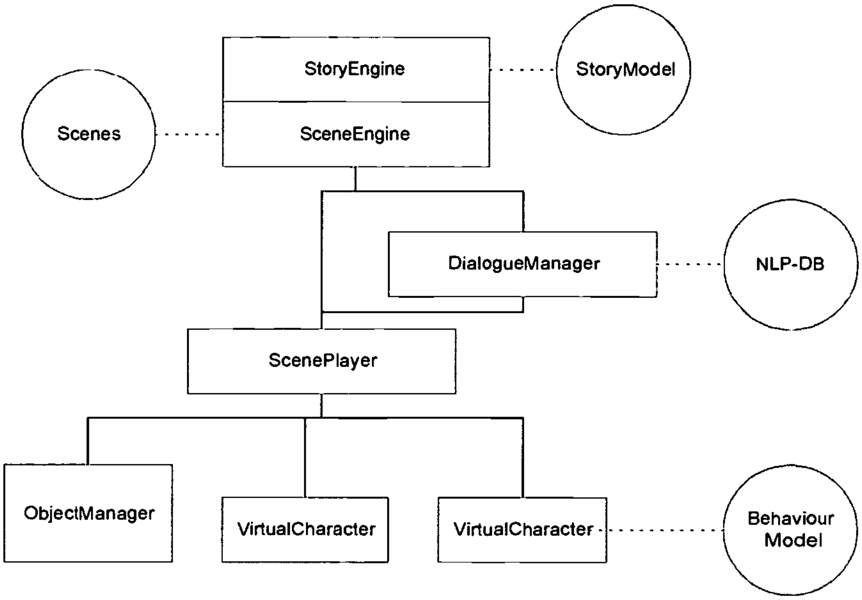

Figure 2: The basic architecture

Story Engine

Our basic approach to the automated generation of stories is to use a central narration component, called "story engine", that combines and adapts prefabricated minimal story components during runtime, according to rules e.g. of narration or pedagogy, and to the interaction results. Those rules are stored as story models. The minimal story components are called "scenes". For details of this process, compare Braun (2002), Spierling et al. (2002). The scene engine plays the scenes that were chosen and adapted by the story engine controlling the subordinate modules.

Dialogue Manager

The dialogue manager is granted control by the story engine if an interactive episode is part of a scene. It is centralized in order to allow using distributed reactions, where, for example, two virtual characters mention two different aspects of a problem as an answer to some question of the user. Its output is a script to the dialogue manager. The script may consist of a single utterance or of a series of commands to the subordinate modules. Currently, the dialogue manager relies on a modified version of the chatbot ALICE (Cf. www.alicebot.org) which makes use of a "natural languagee processing database" (NLP-DB). The use of ALICE is only a provisory solution until some more powerful system is developed.

Scene Player

The scene player executes scene scripts. The scripts control the appearance of the characters and their behaviour at varying levels of abstraction, and govern the objects in the scene; for example, the appearance of icons or the starting of films.

Object Manager

The object manager is controlled directly by the scene engine. It manages the display of the objects of the scene; for instance, the appearance and disappearance of the artworks under discussion.

Virtual Characters

The virtual characters are responsible for the generation of consistent behaviour according to their personality traits and situation appraisal. This behaviour includes turn-taking signals and conversational feedback as well as emotional reactions and goal management. The virtual characters are composed of a mind called agentive person, a simulation of body processes called bodily person, and a graphics output module called manifest person. The parameters of these personality aspects work together to create the final visible personality of the virtual character. The prescriptions may vary in their level of abstraction, ranging from detailed commands of movements to mere annotations in the sentences to say, informing the characters about the pragmatic meaning of the sentences. In the last case, every detail of the performance is in the charge of the virtual characters. Thus, their function is analogous to actors in a theatre, interpreting with leeway the instructions of a director according to their characters and situation awareness. (For more details, cf. Iurgel 2002.)

One of the central research goals of the project is to understand how to combine story elements prepared in advance with less guided, highly interactive episodes. Both aspects are necessary in interactive storytelling, since only the use of prepared dialogues and instructions can guarantee an interesting and consistent storyline, and only the introduction of free interaction enables user adaptation and involvement.

As in any story, the main responsibility for success and suspense lies in the hand of an author. Story creation without an author – "emergent narrative" - is not likely to provide the intended quality of emotional experience. For this reason, in the project "art-e-fact", artists are involved in every step of the development. Thus, suitable authoring tools for interactive storytelling are vital and one of the main challenges of the project (Cf. Scheider, 2003).

Ways of Fostering Art Understading.

Fostering art understanding through the use of a virtual conversation group

This section elaborates more on some aspects of using a virtual conversation group to promote the understanding of art works. In order to concretize the discussion, I introduce two scenarios.

1. Distant horizon scenario. In this scenario, the virtual characters are contemporaries of the human participant and maintain themselves at a distant perspective from the art work, though it has an important meaning for their lives. For example, a scientist who analyses materials and colours, and an historian who knows about cultural context, might be present. This scenario is a good starting point for the participant, because it is easy to understand the virtual characters, and the interest of the characters in the icon can spread to the user.

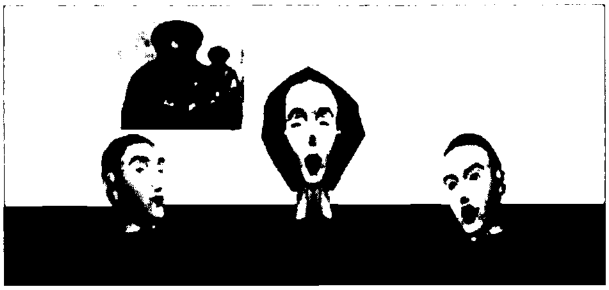

Figure 3: The virtual characters as talking heads

2. Near horizon scenario. Here, the virtual characters live in the past and share the religious fervor that once accompanied the icon. A monk and a mother of a sick child might be present. This constellation must be introduced more carefully, but promises a more intense experience, involving more central feelings like suffering, compassion, hope, and so on.

Advantages of using a conversation group to present the art work on the Web:

1. **Perspective.** Every character can support a definite perspective, e.g. a historian and a scientist. This allows elucidating the theme from different, possibly incommensurable points of view. Personality trait differences have here the important function of establishing a link between opinion and emotion. For example, a reserved and sceptical character is suitable for the role of a scientist, but probably not for the horizons of a religious person.

2. **Dialogic deployment.** The use of a group of characters facilitates the dialogic deployment of thoughts, because insightful sequences of questions, contributions and answers can be prepared in advance. Even if the user is included in the deployment, the presence of a group is technically and conceptually advantageous, because one virtual character can provide the keyword needed by the other to continue the dialogic developments.

3. **Narration.** It would be very difficult to maintain a narrative structure using only a single virtual character because of the unpredictability of user behaviour. If, for example, the author prepares a

narrative structure where the character necessarily has to speak about his early childhood experiences, but the user refuses to keep at it, the story line is likely to fail. This cannot happen if there is another character willing to listen in the place of the user.

4. **Natural language understanding.** Communicating with a single virtual character is difficult simply because it very often won't understand the user. The use of a group alleviates the problem, because it is quite natural that some contributions do not receive a proper answer in a discussion group. Shortcomings of natural language processing do not lead to a breakdown in the conversation.

5. **Social processes.** Social and group processes are among the strongest elicitors of emotional arousal. According to Kemper (1991), there are only two basic dimensions that govern the arousing of emotions in a social relationship; namely, "status" and "power". Those dimensions can be exploited to generate emotionally thrilling plots involving the virtual characters and possibly including the user. For example, the user might be exposed to a planned power change in the course of the narration, his influence as a group member increasing or decreasing depending on his behaviour and on the plot. This can be achieved by a sequence of scenes in which the characters initially disregard the suggestions of the user; in the course of the interaction, his contributions would be "recognized" by the virtual characters as important and increasingly mandatory.

The nowadays prevalent use of anthropomorphic agents on the Web considers only the usage of a single character. I argued that this is a disadvantageous restriction that impedes narration and weakens the interactive experience.

"Interactivity" in interactive storytelling

A non-interactive example for the combination of story and presentation is "Sophie's World" (Gaarder, 1996), where the adventures of a girl serve as a background for explanations in the history of philosophy. Here, the figure of Sophie is used to promote identification, gradual presentation of knowledge, and emotional involvement in the theme.

Certainly, stories in interactive storytelling currently are at risk of not reaching the narrative quality level of non-interactive media where every detail can be controlled by the author or director. Because of this, understanding the role of interactivity in interactive storytelling remains an essential and still not satisfactorily accomplished task. What are the benefits of introducing interactivity into storytelling? Of course, the conflicting goals of the application provide good reasons, even if it should turn out that the intensity of narrative immersion is affected.

Some examples are:

1. **Navigation.** In an edutainment application, it is important that the user can navigate; for example, in order to go back to themes already discussed.

2. **Reviewing and testing.** Interaction is necessary for feedback about some educational goal; e.g. the acquired knowledge could be tested.

3. **Joy of expression.** Expressing personal opinions and observing the virtual character's reactions is a joyful experience.

4. **Exploration.** The user can be engaged in explorative learning strategies, e.g. a virtual character can gradually guide the user to discover more details of an icon.

However, the starting point of this paper was the social and emotional nature of art understanding and the importance of narrative immersion in this context. The question at hand is whether it is possible to use interactivity to intensify the experience of narrative immersion. Traditionally, emotional involvement in a story in achieved by identification with a protagonist. The prospect of placing the user in the protagonist's position in interactive storytelling might be thrilling, but would be opposed to the current goal of rendering another person's emotional and cultural perspective understandable – aside from the technical constraints.

Therefore, I follow the indirect path of creating planned emotional proximity – henceforth "friendship" – between the user and a chosen protagonist virtual character. The reason for this is that friendship is known to boost identification and empathy, and thus, should increase narrative immersion.

An example is a virtual character that is emotionally touched by discovering, in the course of an adventure, the symbolic meaning of a religious icon. Because of the establishment of some sort of friendship with the character, his emotional reaction should also touch the user, promoting his understanding and appreciation of the icon and of its religious context.

Improving Art Understanding Through Friendship With Your Virtual Companion

Previously, I argued that it is possible to foster the narrative immersion through the establishment of some sort of friendship between the user and the virtual character, thus intensifying the identification with the character and his emotions. In pedagogy, those phenomena of learning facilitation under the influence of a liked person fall under the broad and well studied heading of "social learning" (Bandura, 1977, is a classical source). In other words, the idea presented here is that "social learning" is also useful when the main channel is a narration. The research issue under way, now, is to automate the creation of friendship between character and user. In the present context, the relevant psychological factors that determine friendship ("attraction") are the following (Cf. Fehr, 1996, for an overview of those factors):

1. **Similarity.** There is plenty of empirical evidence that similar character traits and opinions are favourable to the establishment of close emotional relationships.

2. **Reciprocity.** Basically, reciprocity means that we tend to like someone who likes us. This dimension is particularly sensitive to temporal factors. For example, increasing sympathy over time is more favourable to the establishment of emotional proximity than is decreasing it, even if the final sympathy level remains the same.

3. **Self-Disclosure.** Self-disclosure is an important indicator of friendship. A person who discloses his feelings to me probably likes me.

4. **Maximized gains in social exchange.** Praise is also very favourable for the establishment of friendship. For example, it is advantageous for me if another person praises my doings, thus raising my status and my self-esteem.

5. **Common enemy.** The existence of a common enemy is also favourable for the establishment of camaraderie as a variety of friendship.

Every one of those factors can be incorporated into a story-telling system. For example, the virtual character chosen for identification can be made to mimic similarity ("You do not like this icon? Nor do I!"); reciprocity ("That's fine that you are enjoying our conversation! I am also learning very much from your contributions!"); or self-disclosure ("I am full of anguish when I think that my answers might not match your questions.") This can already be achieved by integrating an appropriate plot with fairly simple extensions of ALICE, to allow the story engine to guide this chatbot. More powerful methods for interactive episodes, based on the state-update-theory of dialogue moves, are under investigation.

The details and distinctions are tricky and the object of current research. Therefore, the next section will further elucidate the process of automated creation of friendship in interactive storytelling by presenting an example.

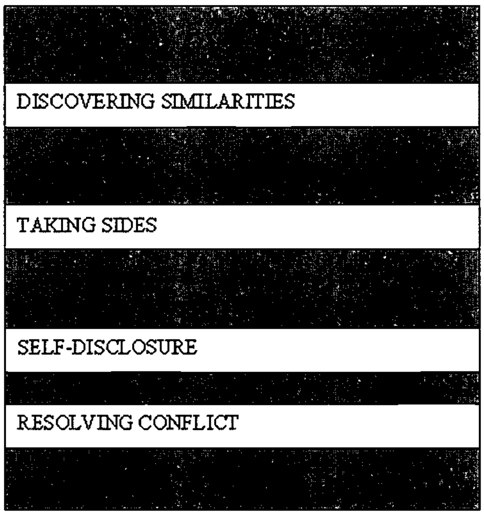

Figure 4: Possible steps of automated friendship forming

Example

In this example, it is assumed that an appropriate story is generated and envelops the scenes under discussion. An abstract initial scheme for the automated establishment of friendship is depicted in figure 4. It is assumed that "discovering similarities", "taking sides", "self-disclosure", and "conflict resolution" are scenes controlled by the story engine according to friendship establishment models.

The creation of this model is a process that involves computer science, psychology, artistic intuition and narrative skills. The scene "taking sides" is included because it allows the virtual character to show how his emotional well-being depends on the behaviour of the user, thus expressing the close friendship he presumes to have with the user. At the same time, the scene "taking sides" creates a common enemy. Short interaction examples follow, in which the system tries to establish friendship between Virtual Character 1 and the user:

DISCOVERING SIMILARITIES

..

Virtual Character 1: "How do you like this icon?"
User: "I love it!"
Virtual Character 1: "I love it too!"
(The virtual character agrees with the user, no matter what he says.)

...

TAKING SIDES

...

Virtual Character 2 (to Virtual Character 1): "You talk too much nonsense!"
Virtual Character 1 (shocked, to User): "Do you think, too, that I talk nonsense?"
User: "Not at all!"
Virtual Character 1 (showing relief and gratitude): "You are a real friend!"
(Virtual Characters 1 has now a slot to show his deep feelings, depending on the side the user takes. Virtual Character 2 becomes antagonist of Virtual Character 1. The scene would have to be extended by the story engine if the user had agreed with character 2, in order to resolve amicably the emerging conflict between user and character 1.)

...

SELF DISCLOSURE

...

Virtual Character 1 (to User. A monk has appeared in the scene.): "I am afraid of this monk. Do you know why?"
User: "No."
Virtual Character 1: "Because he resembles my father!"
(It is important here to take care that the self-disclosure be not ascribed to personality traits of the virtual character, but to his special relationship with the user.)

...

RESOLVING CONFLICT

...

Virtual Character 1 (to Virtual Character 2. Virtual Character 1 is still angry at 2): "Do you still think that I am talking non-sense?!"
Virtual Character 2: "Oh, you are still angry at me. No, please excuse me, I was wrong."
Virtual Character 1: "OK, let us forget about it."

(The conflict was only introduced for the sake of approaching Virtual Character 1 and the user. Since a latent conflict in the group might be harmful to the further development of the story, it is better to resolve it soon.)

...

Conclusion

I presented the project Art-E-Fact and the use of a narrative, interactive discussion group to deepen and contextualize art presented on the Web. I referred to art theories which state that art understanding is an inherently social, emotional and cultural process, and that therefore narration and dialogue are appropriate means of fostering art experience. I presented an overall initial view of the system and argued that creating friendship with the virtual companions is likely to improve the process of identification and narrative immersion, and sketched the automated processes that establish the desired emotional proximity.

The ongoing issue might be regarded as a first attempt to virtually populate spaces that were until now wasteland like the Web, providing moments of automated suspense, immersion, and joyful learning to the interested user.

References

André, E., M, Klesen, P. Gebhard, S. Allen, & Th.Rist (2000). Exploiting Models of Personality and Emotions to Control the Behavior of Animated Interface Agents. Jeff Rickel (eds.), *Proceedings of the workshop on "Achieving Human-Like Behavior in Interactive Animated Agents"*, in conjunction with the Fourth International Conference on Autonomous Agents, pp. 3-7, Barcelona, June 2000.

Bandura, A. (1997). Social Learning Theory, Prentice-Hall, Englewood Cliffs, N.J. 1977.

Bearman, M. (2001). Narrative and Cases: Implications for Computer-based Education. *ASCILITATE*, December 7-10.

Braun, N. (2002). Automated Narration – the Path to Interactive Storytelling. *NILE 2002*, Edinburgh, Scotland.

Fehr, B. (1996). *Friendship processes*. Sage Publications, California.

Gadamer, H.-G. (1993). *Truth and Method*. Continuum Publisher.

Gaarder, J. (1996). *Sophie's World: A Novel About the History of Philosophy*. Paulette Moller (Translator). Mass Market Paperback.

Graesser, A.C., & V. Ottati (1995). Why Stories? Some Evidence, Questions, and Challenges. In R. S. Wyer (Ed.), *Knowledge and Memory: The Real Story*. (pp. 121-132). New Jersey: Lawrence Erlbaum Associates.

Heidegger, M. (1973). The Origin of the Artwork. In: *Philosophies of Art and Beauty* (...). Albert Hofstadter and Richard Kuhns (eds.). University of Chicago Press.

Iurgel, I. (2002) Emotional Interaction in a Hybrid Conversation Group. *PRICAI-02. Workshop on Lifelike Animated Agents*. Tokyo, Japan. JSPS.

Kemper, Th. D. (1991). Predicting Emotions from Social Relations. *Social Psychology Quarterly*, Vol. 54, No. 4, 330-342.

Machado, I, Paiva, A. & Brna, P. (2001). "Real Characters in Virtual Stories", to appear in the "1st International Conference on Virtual Storytelling, *Springer.*

Mateas, M. & A. Stern. *(2000). "Towards integrating plot and character for interactive drama"*, Proceedings AAAI Fall Symposium on Socially Intelligent Agents: The Human in the Loop.

Paiva, A. & I. Machado. (2001). "Life-Long Training with Vincent, a Web-based Pedagogical Agent", in the *International Journal of Continuing Engineering Education and Life-Long Learning*, Spring Issue.

Rickel J,, S. Marsella, J. Gratch, R. Hill, D. Traum, and W. Swartout (2002). *Toward a New Generation ofVirtual Humans for Interactive Experiences.* IEEE Intelligent Systems 17(4), July/August, pp. 32-38. (Special issue on AI in Interactive Entertainment.)

Schank, R.C., & R.P. Abelson (1995). Knowledge and Memory: The Real Story. In R. S. Wyer (Ed.), *Knowledge and Memory: The Real Story.* (pp. 1-86). New Jersey: Lawrence Erlbaum Associates.

Schneider, O. (2003). *Storyworld Creation: Authoring for Interactive Storytelling.* WSCG (to appear).

Spierling, U.; Braun, N.; Iurgel, I.; Grasbon, D. (2002). Setting the scene: playing digital director in interactive storytelling and creation. *Computers & Graphics 26*, 31-44.

Delivering Digital Heritage

Using Cinematic Techniques in a Multimedia Museum Guide

M. Zancanaro, O. Stock, I. Alfaro, ITC-irst Italy

Abstract

In this paper we introduce the idea of enhancing the audio presentation of a multimedia museum guide by using the PDA screen to travel throughout a fresco and identify the various details in it. During the presentation, a sequence of pictures is synchronized with the audio commentary, and the transitions among the pictures are planned according to cinematic techniques. The theoretical background is presented, discussing the language of cinematography and the Rhetorical Structure Theory to analyze dependency relationships inside a text. In building the video clips, a set of strategies similar to those used in documentaries was employed. Two broad classes of strategies have been identified. The first class encompasses constraints imposed by the grammar of cinematography, while the second deals with conventions normally used in guiding camera movements in the production of documentaries. The results of a preliminary evaluation are also presented and discussed.

Keywords: Multimedia Museum Guides, Cinematography, Interaction on Small Devices, Location-awareness

1. Introduction

Many research projects are exploring the new possibilities offered by Personal Digital Assistants (PDAs) in a museum setting (for example, see Grinter et al, 2002, Cheverst 2000 and Not et al., 1998). Usually, these multimedia guides use static images, while others employ pre-recorded short video clips about museum exhibits. In a previous work (Not et al, 1998, 2000), we explored different techniques to automatically build multimedia, location-aware presentations in a museum setting. The advent of more powerful devices has allowed researchers to experiment with new forms of multimedia, in particular time-based media such as animations.

In this paper we introduce the idea of enhancing the audio presentation (dynamically assembled prerecorded or synthesized speech) of a complex fresco by using the PDA screen to travel throughout the fresco itself and identify details. At presentation time, a sequence of pictures is synchronized with the audio commentary, and the transitions among them are planned according to cinematic techniques. Our hypothesis is that the use of this type of animation to present the description of a painting allows the visitor to better identify the details introduced by the audio counterpart of the presentation. In this manner, both the efficiency and the satisfaction dimensions of the system usability are increased (Nielsen, 1994) while also providing an enhanced learning experience for the visitor.

The language of cinematography (Metz, 1974), including shot segmentation, camera movements and transition effects, is employed in order to plan the animation and to synchronize the visual and the verbal parts of the presentation. In building the animations, a set of strategies similar to those used in documentaries was thus employed. Two broad classes of strategies have been identified. The first class encompasses constraints imposed by the grammar of cinematography, while the second deals with conventions normally used in guiding camera movements in the production of documentaries. For instance, a strategy in the first class would discourage a zoom-in immediately followed by a zoom-out, while a different strategy in the second class would recommend the use of sequential scene cuts, rather than a fade-out effect, to visually enumerate different characters in a scene. It is worth noting that in the latter strategy it is often necessary to make reference to the discourse structure of the audio part of the presentation, such as enumeration of properties, background knowledge, and elaboration of related information. In order to formally use discourse structure, we employ the Rhetorical Structure Theory (Mann and Thompson, 1987).

At present, we have completed a first prototype of a multimedia guide that employs cinematic techniques in presenting information for a fresco at Torre Aquila in Trento, Italy. A Web-based simulation of the multimedia guide can be seen at http://peach.itc.it/preview.html.

The next section briefly discusses the issues in designing a multimedia museum guide. Section 3 introduces the features of the Torre Aquila prototype. Sections 4 and 5 present the theoretical background, discussing respectively the relevant concepts for the language of cinematography and the Rhetorical Structure Theory to analyze dependency relationships inside a text. Section 6 illustrates the strategies used in our multimedia guide to produce effective and pleasant video clips starting from audio commentaries. Finally, in section 7, the results of a preliminary evaluation are presented and discussed.

2. The Museum as a Smart Environment

Museums and cultural heritage institutions recreate an environment conducive to exploring not only the exhibited objects and works of art, but also new ideas and experiences. Visitors are free to move around and learn concepts, inquire and even apply what is leaned to their own worldview. A museum visit is thus a personal experience encompassing both cognitive aspects, such as the elaboration of background and new knowledge, and emotional aspects that may include the satisfaction of interests or the fascination with the exhibit itself. Despite the inherently stimulating environment created by cultural heritage institutions, on their own they often fall short of successfully supporting conceptual learning, inquiry-skill-building, analytic experiences or follow-up activities at home or the school (Semper and Spasojevic, 2002).

The optimal multimedia tourist guide should support strong personalization of all the information provided in a museum in an effort to ensure that each visitor be allowed to accommodate and interpret the visit according to his own pace and interests. Simultaneously, a museum guide should also provide the appropriate amount of impetus to foster learning and self-development so as to create a richer and more meaningful experience.

In order to achieve the above goals it is necessary that the information be presented in a manner that is appropriate to the physical location of the visitor as well as to the location of the works of art within the environment. Smoothly connecting the information found in an exhibit and presenting it to the visitor in a flexible yet coherent manner with respect to his physical location can mazimize the overall experience and absorption of new information for the viewer (Stock and Zancanaro, 2002). In other words, if the information is provided in a manner that flows and relates pieces to each other, this process in and of itself can aid in stimulating the visitor's interest and, hence, desire to inquire, analyze and learn. This idea relates to the concept of *situation-aware content*, where information is most effective if presented in a cohesive way, building on previously delivered information. This may be accomplished by using comparisons and references to space and time, which in turn may aid the visitor in becoming oriented within the museum as well as across the various works of art.

The ideal audio guide should not only guess what the visitors are interested in, but also take into consideration what they have to learn: orienting visitors, providing opportunities for reflection and allowing them to explore related ideas, thereby greatly enhancing the visit's educational value. In essence, the guide should stimulate new interests and suggest new paths for exploring the museum. A system that supports visitors in their visit should take into account their agenda, expectations and interests as well as the peculiarities of a cultural experience in a physical environment.

It is essential to also consider the importance of creating an overall experience that truly addresses the needs of a person visiting a museum. This requires not only providing the visitor with a vast amount of information, even if wonderfully presented, but also allowing the person to spend a pleasurable and entertaining time at the exhibit. The concept of the *immersive environment* addresses the importance of creating a technology that supports rather then overwhelms the real experience of visiting a museum. A museum guide of this nature must be able to create a balance in terms of attention required from the visitor, also allowing time to be spent enjoying the "romance" of the cultural heritage institution and the works found therein.

These and other challenges come into play when designing a system for the entertainment and edutainment of museum visitors. Creating an elec-

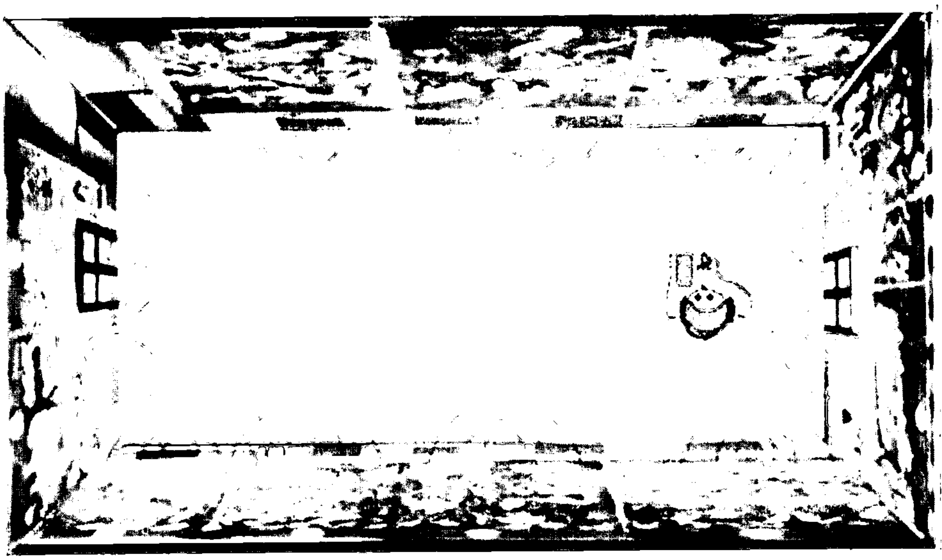

Figure 1: Torre Aquila and the grid of infrared

tronic tourist guide that transforms the user experience from one of simple consultation (commonly achieved with audio guides, multimedia kiosks, CDROMs or even books) to an immersion into a rich information environment indeed requires a careful examination of all the abovementioned factors, while also considering input from visitors themselves. Difficulties arise when observing that such systems are not intended to help users perform specific work-related tasks, and most of the time they cannot be brought back to clearly stated user requirements. Keeping in mind that the ultimate goal of a guide is to engage the user and to stimulate learning, it becomes clear that the nature of this kind of system imposes a balance between the designer's vision and user needs.

Using animations or video clips enhance the richness of the interaction though these may also distract the user by calling attention to the device rather than to the exhibit itself. Our hypothesis is that, on the contrary, a carefully planned video clip describing the exhibit will actually help the visitor quickly localize the details of the painting as well as aid the flow of the presentation by illustrating the relationship between new and already presented information.

3. The Prototype at Torre Aquila

We have applied the idea of using cinematic techniques for presenting details of artworks in a prototype of a multimedia guide for Torre Aquila' a tower at the Buonconsiglio Castle in Trento, where a fresco called "The Cycle of the Months", a masterpiece of the gothic period, is found. This fresco, painted in the Fifteenth Century, illustrates the activities of aristocrats and peasants throughout the year. The fresco is composed of eleven panels, each one representing one month (the month of March was destroyed over time) and occupies the four walls of the tower (see figure 1).

Our multimedia guide, implemented with Macromedia Flash on a PDA, detects the position of the visitor by means of infrared emitters placed in front of each panel. Interaction with the system is both proposed by the system itself and accepted by the user, thus sharing the responsibility of information access. When the system detects that the visitor is in front of one of the four walls, a picture of that wall is displayed on the PDA and, after a few seconds, if the user has not changed position, that panel is highlighted (see figure 2). At this point, the visitor can click on the panel and receive a multimedia presentation of the panel chosen.

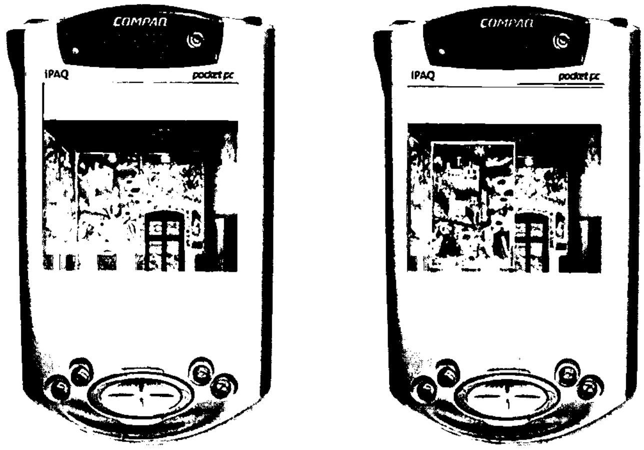

Figure 2: Snapshots of the multimedia guide localizing the user

The multimedia presentation is composed of an audio commentary accompanied by a sequence of images that appear on the PDA display and help the visitor quickly identify the fresco's details mentioned in the commentary. For instance, when a specific detail of the panel is explained by the audio, the PDA may display or highlight that detail, thus quickly calling the attention of the user to the area in question.

During the presentation, the PDA displays a VCR-style control panel and a slide bar to signal the length of the video clip and its actual position (see figure 3). At any given moment, the user is free to pause, fast forward, rewind and even stop the presentation by tapping on the appropriate control panel button. In this manner, the visitor is able to control the speed as well as the information itself, while also revisiting sections found most interesting.

Figure 3: Snapshot of the multimedia guide playing a video clip.

4 .The Language of Cinemaatography

According to Metz {1974), cinematic representation is not like a human language that is defined by a set of grammatical rules; it is nevertheless guided by a set of generally accepted conventions. These guidelines may be used for developing multimedia presentations that can be best perceived by the viewer. In the following, we briefly summarize the basic terminology of cinematography. In section 6 we will discuss how these conventions can be expressed both in terms of constraints on camera movements and in terms of strategies related to the discourse structure of the associated audio commentary.

4.1 Shot and camera movements

The shot is the basic unit of a video sequence. In the field of cinematography a shot is defined as a continuous view from a single camera without interruption. Since we only deal with still images, we define a shot as a *sequence of camera movements applied to the same image*.

The basic camera movements are *pan*, from "panorama", a rotation of the camera along the x-axis, *tilt*, a rotation along the y-axis, and *dolly*, a rotation along the z-axis.

4.2 Transition effects

Transitions among shots are considered the punctuation symbols of cinematography; they affect the rhythm of the discourse and the message conveyed by the video.

The main transitions used are *cut*, *fade*, and *cross fade*. A cut occurs when the last frame of a shot is immediately replaced by the first frame of the following shot. A fade occurs when one shot gradually replaces another one, either by disappearing (fade out) or by being replaced by the new shot (fade in). A particular case of a fade happens when instead of two shots, there is one shot and a black screen that can be, again, faded in or faded out. Finally, a cross fade (also called dissolve) occurs when two shots are gradually superimposed during the moment when one is faded out while the other is faded in.

5. Rhetorical Structure Theory

Rhetorical Structure Theory (Mann and Thompson, 1987) analyses discourse structure in terms of dependency trees, with each node of the tree being a segment of text. Each branch of the tree represents the relationship between two nodes, where one node is called the nucleus and the other is called the satellite. The information in the satellite relates to that found in the nucleus in that it expresses an idea related to what was said in the nucleus. This rhetorical relation specifies the coherence relation that exists between the two portions of text contained in the nodes. For example, a *Cause* rhetorical relation holds when the satellite describes the event that caused what is contained in the nucleus. Figure 4 shows an example of a rhetorical tree. Here the second paragraph provides background information with respect to the content expressed in the first paragraph. This additional information acts as a sort of reinforcement for what was previously said in the first paragraph and consequently facilitates the absorption of information. In the original formulation by Mann and Thompson, the theory posited twenty different rhetorical relations between a satellite and a nucleus, while other scholars have since added to this theory.

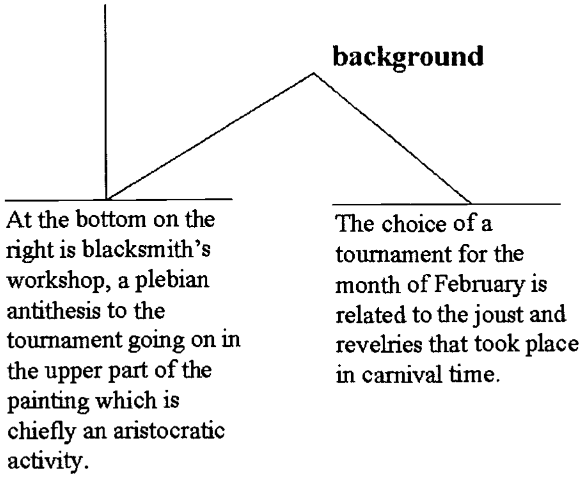

Figure 4: An example of a rhetorical tree (simplified).

RST was originally developed as part of work carried out in the computer-based text generation field. In a previous work (Not and Zancanaro, 2001), we described a set of techniques to dynamically compose adaptive presentations of artworks from a repository of multimedia data annotated with rhetorical relations. These techniques have been exploited in an audio-based, location-aware adaptive audio guide described in Not et al., {2000). The audio commentaries produced by this audio guide are automatically annotated with the rhetorical structure. In the next section we will discuss how this information can be used to create more effective video clips to accompany the commentary.

6. Video Clips on Still Images

Video clips are built by first searching for the sequence of details mentioned in the audio commentary, deciding the segmentation in shots, and then planning the camera movements in order to smoothly focus on each detail in synchrony with the verbal part.

In building a video clip, a set of strategies similar to those used in documentaries is employed. Two broad classes of strategies have been identified. The first class encompasses constraints imposed by the grammar of cinematography, while the second deals

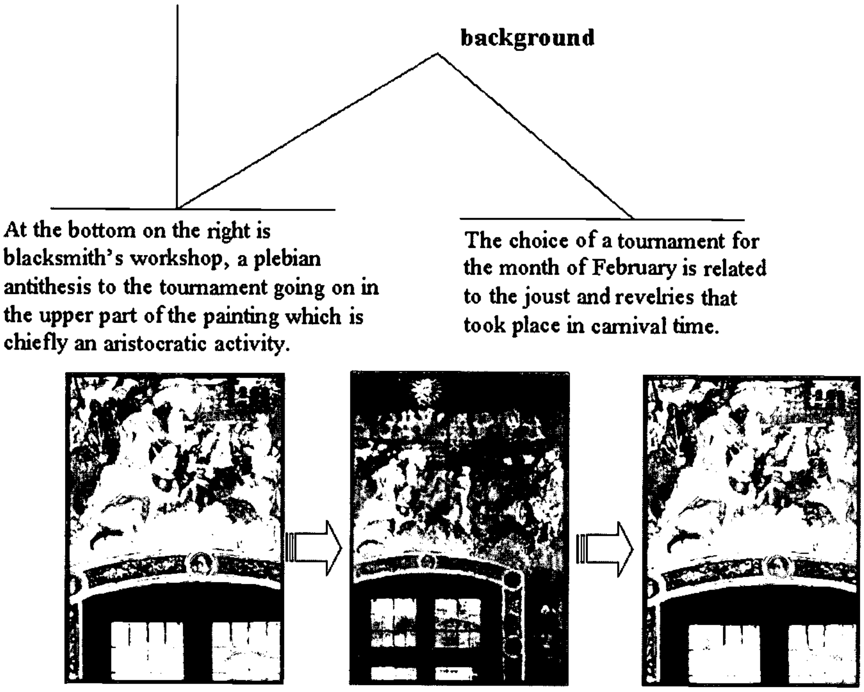

Figure 5: The "Tournament" example: from the text to the video clip.

with conventions normally used in guiding camera movements in the production of documentaries.

While the constraints are just sequence of forbidden camera movements, the conventions are expressed in terms of rhetorical structures found in the audio commentary. In our view, the verbal part of the documentary always drives the visual part.

6.1 Constraints on Camera Movements

In order to ensure a pleasant presentation, constraints on camera movements have to be imposed. For example, a pan from right to left forbids a subsequent pan from left to right. In general, applying any given movement (pan, tilt and zoom) and then immediately reapplying it on the reverse direction is discouraged because this action renders the video uncomfortable to watch.

Given that the audio commentary drives the visual part, it is often the case that such forbidden combinations of camera movements are required. In these cases, two tricks can be applied: either choosing a different way of focusing the detail required by the verbal part; for example a zoom out can often effectively replace a pan, or starting a new shot altogether. In the latter case, the two shots should be linked by a transition effect that suggests continuity, such as a short fade.

6.2 Rhetorical Strategies

Constraints on camera movements alone are sufficient to ensure a pleasant presentation, yet they do not impact the effectiveness of the video clip. In order to have a more engaging presentation, the visual part should not only focus on the right detail at the right time, but also support the presentation of new audio information by illustrating its relation to information that has been already given. In this manner, continuity between the pieces of information is built, and in turn facilitates the viewing of the video clip while stimulating the absorption of new information.

The text in figure 5 can be visually represented with two shots of the same image (that is, the tournament) linked by a long cross fade. Technically, having two shots is not necessary, since the image is the same, but the cross fade helps the user understand that background information is going to be provided. The first image is thus presented while the first paragraph is heard over the audio; then when the audio switches to, in this case, the background information, the image is enlarged to cover the entire panel and finally is refocused on the detail once the audio has stopped.

A rhetorical strategy suggests, on the basis of a rhetorical tree configuration, what shot segmentation and which transition effect should be applied. The strategies employed in the Torre Aquila multimedia guide were elicited by a focus group activity with a documentary director.

7. Preliminary Evaluation

A formal evaluation of the prototype will start next March at Torre Aquila. Preliminary studies and pilot tests show encouraging results and interesting effects.

All users became acquainted with the system very quickly. Most of them used the PDA as a "3D mouse", pointing directly to the infrared emitters to speed up the localization. Future investigations will evaluate how users can be more directly involved in the process of localization.

Most of the users complained before actually using the system that a video sequence on a PDA would distract their attention from the real artwork. After a short interaction with the system, however, they appreciated the possibility of quickly localizing small details on the fresco. This demonstrates that use of cinematic techniques in a multimedia guide can be effective, particularly in explaining complex painting. The different effects that the verbal and the visual parts of the presentation have on the user's attention are yet to be investigated.

8. Conclusion

This paper discussed how cinematic techniques can be used in a multimedia museum guide to provide more pleasant and effective presentation of information. Video clips are built by first searching for the sequence of details mentioned in the audio commentary, deciding the segmentation in shots, and then planning the camera movements soas to smoothly focus on each detail in synchrony with the verbal counterpart. In our approach, the verbal part always drives the visual part.

The video clips are built accordingly to two broad classes of strategies. The first class encompasses constraints imposed by the grammar of cinematography, while the second deals with conventions normally used in guiding camera movements in the production of documentaries.

While the constraints are just a sequence of forbidden camera movements, the conventions are expressed in terms of rhetorical structures found in the audio commentary. By coupling these cinematic techniques into organized guidelines, the creation of multimedia video clips can vastly help to improve quality as well as the effectiveness of the presentations. A visitor to a museum can thus benefit from an automatic guide that causes minimal interference with the enjoyment and learning experience provided by an exhibit. As a case study, a multimedia museum guide for Torre Aquila in Trento has been presented and the results of a preliminary evaluation have been discussed

Acknowledgments

This work has been supported by the PEACH and TICCA projects, funded by the Autonomous Province of Trento.

References

Metz 1974. *Film Language: a Semiotics of the Cinema.* Oxford University Press, New York.

Cheverst, K., N. Davies, K. Mitchell, A. Friday and C. Efstratiou, 2000. Developing a Context-aware Electronic Tourist Guide: Some Issues and Experiences *Proceedings ofCHI 2000.* Amsterdam.

Grinter, R.E., P. M. Aoki, A. Hurst, M. H. Szymanski, J. D. Thornton and A. Woodruff, 2002. Revisiting the Visit: Understanding How Technology Can Shape the Museum Visit. In *Proc. ACM Conf. on Computer Supported Cooperative Work,* New Orleans, LA.

Mann, W.C. and S. Thompson, 1987. Rhetorical Structure Theory: A Theory of Text Organization, In L. Polanyi (ed.), *The Structure of Discourse*, Ablex Publishing Corporation.

Nielsen, J. 1994. *Usability Engineering*, Morgan Kaufmann, San Francisco.

Not, E., D. Petrelli, O. Stock, C. Strapparava and M. Zancanaro, 2000. The Environment as a Medium: Location-aware Generation for Cultural Visitors. In Proceedings of the workshop on *Coherence in Generated Multimedia*, held in conjunction with INLG'2000, Mitze Ramon, Israel.

Not, E., D. Petrelli, M. Sarini, O. Stock, C. Strapparava, M. Zancanaro, 1998. *Hypernavigation in the Physical Space: Adapting Presentations to the User and the Situational Context.* In New Review of Hypermedia and Multimedia, vol. 4.

Not and Zancanaro, 2001. E. Not, M. Zancanaro. "Building Adaptive Information Presentations from Existing Information Repositories". In Bernsen N.O., Stock O. (eds.), Proceedings of the International Workshop on Information Presentation and Multimodal Dialogue, Verona, Italy.

Semper, R., and M. Spasojevic, 2002. The Electronic Guidebook: Using Portable Devices and a Wireless Web-Based Network to Extend the Museum Experience. In Proceedings of Museums and the Web Conference, Boston, MA.

Stock, O., M. Zancanaro. Intelligent Interactive Information Presentation for Cultural Tourism. Invited talk at the *International Workshop on Natural, Intelligent and Effective Interaction in Multimodal Dialogue Systems.* Copenhagen, Denmark. June, 2002.

The Use of an Information Brokering Tool in an Electronic Museum Environment

Andreas Zimmermann, Andreas Lorenz, and Marcus Specht, Fraunhofer Institut for Applied Information Technology, Germany

Abstract

When art and technology meet, a huge information flow has to be managed. In the LISTEN project conducted by the Fraunhofer Institut in St.Augustin, we augment every day environments with audio information. In order to distribute and administer this information in an efficient way, we decided to employ an information brokering tool for the management of information items. Furthermore, the generation of user profiles and the personalized presentation of information are possible by this means. This contribution depicts an approach of transferring our information brokering experience to Web museum applications. We show how the LISTEN domain model can easily be extended by an overlay model and adapted to a Web museum environment.

Keywords: Audio-Augmented Environments, Information Brokering, Ontology Modeling, User Modeling, Personalization, Context-Awareness

Introduction

A variety of guidance systems or information systems have been developed in the last few years for the support of a museum visit or for preparation for such a visit. In most cases the authenticity and the possibility of contextualizing the information presentation to the current position or situation of a user were seen as central issues. Furthermore, classical museum audio guides were much more acceptable to museum visitors because of the easy handling and the quality of sound presentation. For the flexible use of information about art objects and museum exhibitions, we propose a centralized model around a domain ontology that is described in a software tool for information brokering. We think that the proposed model is a very flexible way to reuse existing information and support curators and exhibition experts to in the design of a variety of personalized museum experiences, ranging from an exhibition Web site to a interactive audio experience in the museum space to a personalized CD-ROM production for taking home.

Context as a mean for adaptation of information selection and presentation has been described in a variety of ways and approaches (Gross & Specht, 2001; Shilit, Adams & Want, 1994; Dey & Abowd, 1999). Nevertheless the underlying problem of identifying similarities and differences between various constellations of context parameters has not been discussed intensively in the literature. From our point of view identifying important context parameters to describe user interaction is an essential issue when designing context aware information services.

The identification and the description of the context parameters is a non-trivial task. A central issue is to find a solution for structuring an information domain appropriate not only from the point of view of an information engineer but also from the point of view of a user. Especially for designing personalized information services the structure and the intuitiveness of the information structuring is essential for the successful application of user modeling and personalization methods.

The LISTEN project conducted by the Fraunhofer Institut in St.Augustin deals with the audio augmentation of real and virtual environments. A first LISTEN prototype has been successfully installed at the Kunstmuseum in Bonn (Unnützer, 2001). Visitors walk through the museum and experience personalized audio information about paintings by August Macke through wireless headphones. In such a scenario, the distribution and authoring of this valuable information on exhibits is a non-trivial task.

Audio-Augmenting Real and Virtual Environments

Combining high-definition spatial audio rendering technology with advanced user modeling methods creates audio-augmented environments. Visitors are immersed in a dynamic virtual auditory scene that consistently augments the real space they are exploring. The physical environment is augmented through a dynamic soundscape, which users experience over motion-tracked wireless headphones for 3D spatial reproduction of the virtual auditory scene.

A sophisticated auditory rendering process takes into account the current position and orientation of the visitor's head in order to seamlessly integrate the virtual scene with the real one. Speech, music and sound effects are dynamically arranged to form an individualized and situated soundscape offering exhibit-related information as well as creating context-specific atmospheres. Next to the automatic adaptation of sound scene rendering to the position and orientation of the user's head, the audio stream is controlled in two ways: events (mediated interaction) that are used to start and stop the playback of information items in form of audio recordings; and continuous control (immediate interaction) changing parameters in the audio-generation of the presentation (e.g. a sound that gets continuously louder as you approach a certain position within the space).

The dynamic composition of the soundscape is personalized through each visitor's spatial behavior, the history of the visit, and interests or preferences either expressed explicitly by the visitor or inferred from the visitor's behavior (Eckel, 2001). To present individualized media and create augmented environments, such systems have several models in common (Goßmann & Specht, 2001) that are described in more detail in this section: the World model, the augmentation layer, the domain model and the user model.

The *World Model* (*Space Model, Location Model*) describes the physical environment the user moves through while interacting with the system. In the LISTEN environment, the space model contains the geometric information of the exhibition space and its objects. The LISTEN world model is a detailed VR-based geometric model. It is created for the AVANGO application (Tramberend, 1999) and is described as a geometric scene graph. Therefore, a LISTEN environment can be tested and prototyped in a CAVE system (Eckel, 2001), or be explored in real space with virtual audio content displayed through a wireless motion-tracked headphone.

The *Augmentation Layer* on top of the World Model defines areas (Zones, Segments, Triggers) within the world model that contain active information or sound objects the users of the system interact with. The augmentation layer filters the position and motion of the user by dividing the dimensions the user moves through (location and orientation) into meaningful constraints and deriving continuous parameters from them. By defining zones and segments, the visitor's focus obtains a valuable meaning.

The *Domain Model* holds information about sound objects and other hypermedia objects connected to the physical space via the augmentation layer by using meta data. The domain model builds up a virtual acoustic space in which the location of virtual sound sources and spaces are defined. Stopping in front of an exhibit generates aural information about the art piece. Moving the head and body activates a further audio source, where music deepens the visitor's impressions, or the voice of a commentator talks about the artist or describes the period the painting originates from.

The *User Model* contains knowledge and profile information about the system's users. While the user moves in physical space, events are sent to the user model, and by these events the model is refined. The user model stores knowledge about the visitor - like preferences, interests in arts, and the history of the visit. In combination with the visitor's spatial position (delivered by the tracking system), the content of the user model strongly influences the presentation of information according to the current visitor's context.

Information Brokering

In this real world application of LISTEN, we decided to use an information brokering tool for authoring valuable information on exhibits and distributing requested information among the visitors in an efficient way. Information brokering is a value adding process of mediation between information demands and information offers. The added value emerges from the understanding of the domain complexity and the definition of a useful vocabulary. Additional structures and interpretation rules (implicit and explicit) in information exchange simplify the processing and comprehension of exchanged information. Figure 1 illustrates the interrelation between data, information, and knowledge.

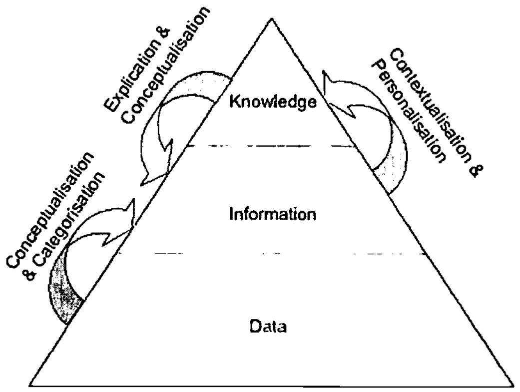

Figure 1

Three roles participate in the information brokering process: the provider offers information, the consumer demands information, and the *broker* mediates between the other two. The quality of the information brokering process depends to a great extent on the knowledge available to the broker. Knowledge about available sources, the domain, consumers and other brokers is needed. *Source knowledge* is created in the domain representation or maintenance processes and describes the quality of sources and how they can be accessed. *Domain Knowledge* is about the contents of the brokering domain and should reflect the provider's understanding of the domain as well as the consumers' perception of the domain. *Consumer Knowledge* is created in the consumer-oriented process and describes the consumer and his concrete information need. Consumer knowledge has to map onto domain knowledge to fulfill the information need. To ensure an optimal service to consumers, they should be served by the best broker according to their information need. This assignment task depends on the availability of *Expert Knowledge* about different brokers.

In the LISTEN project, the brokering tool can be understood as a server mediating information items between the exhibition's curator and the LISTEN application. In the role of the information provider, the curator defines the domain model and authors the information items, as illustrated in the left column of Figure 2. From the point of view of the information broker, the information is consumed by the LISTEN application, although it finally passes the information to the visitor as the real consumer. In the role of the information consumer (Figure 2, right column), the LISTEN application defines an interest model regarding to the visitor's location, preferences, interests, and history of already visited objects.

To mediate between the provider and the consumer, a brokerage tool acts as a server to deliver filtered information on demand. Such a brokerage development and management environment (the *Broker's Lounge*; Jarke, Klemke & Nick, 2001) has been developed at the Fraunhofer Institut for Applied Information Technologies (FIT) in St. Augustin. With the aid of this application, a large variety of scenarios within the general framework of Figure 1 and Figure 2 can be quickly developed, and efficiently and flexibly executed.

Information brokering processes make use of and create a number of information items that describe single units of information (cf. Klemke & Koenemann, 1999). As Klemke (2002) has defined, each information item is an instantiation of a concept which describes the structure of the brokered items. In order to organize information items, categories describe fundamental principles or ideas. In many cases, these categories define hierarchical trees. In Figure 1 *conceptualization and categorization* is illustrated as a procedure for the concretion of raw data. The next section provides an outline of definition of the LISTEN project's information items and illustrates a hierarchical tree to categorize them.

Modeling Information Items

Bringing arts and technology together means the administration and distribution of a huge amount of information. In this paper we propose the use of

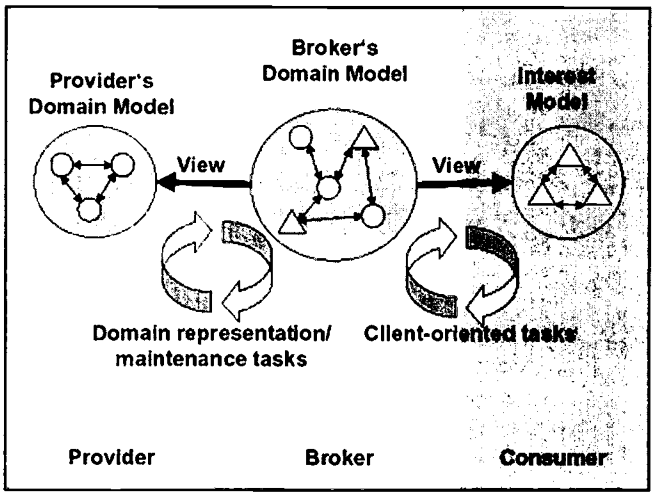

Figure 2

BEST COPY AVAILABLE

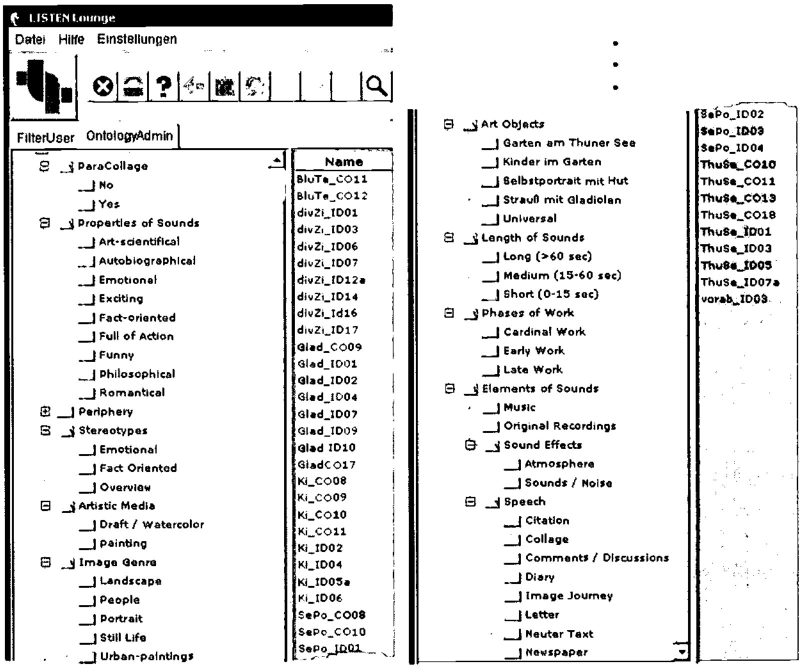

Figure 3

a central management component for all information items. We present an approach of modeling information items in a museums environment by means of domain ontology in an information brokering tool. In the case of the LISTEN project, the set of information items is composed of sound items. We describe the application of the ontology methodology (*concepts* and *categories*) to generate a meaningful domain model and depict the association of information items to exhibition objects.

In recent projects developing guiding systems and electronic art guides, the structuring and internal representation of information appeared to be a central issue for delivering the right information at the right time to a specific user. Even more, for the personalization of presentations and the contextualized delivery of information pieces, the underlying structure of the represented information needed to be a core part of the work. Nevertheless there was a trade off between the efforts of authoring information items into a highly enriched information representation and the daily work of curators and information providers for museum environments.

To find the right balance between user efforts for authoring and creating museum information and meta data in previous projects, we often used an object-oriented approach, where the main entities were the art objects as such. This led us to typical presentation and structuring of information as found in art databases and museum information systems.

By detailed user needs analysis of museum visitors and by the analysis of human guides, we realized that the event character of museum visits in most cases is much more important than plain information about the art object as such. The presentation of data about artworks is just one style of presentation that can be appropriate but in practice is rarely used by human guides. Furthermore, through our work with museum curators and artists in the LISTEN workshops we elicited different strategies and methodologies to structure information items and to combine those information items in a highly flexible way for different users.

The main entities for generating presentations are therefore not the artworks as such, but the parts or chunks of a presentation that can come from a diversity of sources and use a variety of stylistic means to make the museum visit an interactive experience. Therefore, for creating interactive audio augmented LISTEN spaces, we decided to choose the sound item as the main entity. Sound items can be classified in a category system with several dimensions describing the sound items technically and stylistically. An overview of the dimensions can be seen in Figure 3. This figure illustrates the domain model inside the graphical interface of the information brokering tool, organized as a tree, and a list of information items of the LISTEN project.

The single sound items as such are independent episodes or chunks of presentation that can be combined flexibly because they contribute to a variety of art objects. The expressive power of such a structure can be easily explained with an example:

In an introduction to several artworks of August Macke, the different artworks are connected by personal episodes, events and experiences of August Macke, by reactions and interaction with his social environment (personal letters, media, press articles), or by the "Zeitgeist" and other factors of the environment of August Macke at that time. In most cases human guides extend the visual perception of such artworks with these "stories or impressions" of that time period to immerse the user in an authentic experience of the works and the life of the artist. By the structured description (meta tagging) of a huge collection of such "stories or impressions," the LISTEN system allows a highly flexible way of "immersing the user" into the August Macke experience.

Therefore we designed the domain ontology (meta data) for the Macke-Exhibition to allow for the description of small information items on a variety of dimensions, allowing for the connection and the individualized sequencing and presentation of these items. Beside the simple classification of the information items in a tree structure, the information brokering tool allows for classification of sound items into multiple categories.

The domain ontology contains:

- technical descriptions of the sound items; such as length of item, type (music, speech, sound effects)
- classification concerning the relation to the physical space objects (art objects to which an item contributes, physical area zones or focuses to which they are connected)
- classification about phases of work, image genre, or art technical aspects
- classification about the preferred target group, like the stereotypical listeners for such a sound item, or the emotional impacts or dramaturgy.

Especially speech sound items could be further classified into subcategories like Citation, Collage, Diary, Letter, Newspaper and others to describe their style of presentation. As described in the example of the Macke-Exhibition, the multidimensional classification of the sound items allows for a variety of sequences and presentation styles, even combining sound items on several channels; i.e. typically music, effects, and speech. An enormous advantage of the description in such away is that the curator of an interactive experience is not forced to design a complete sequence of information presentation but can combine resources in a collage style or define sequencing rules on the level of possible connection categories and not for single sound items.

In the LISTEN application, the visitor's position and orientation are observed by a tracking system and interpreted within the context of a virtual environment, which connects the real world objects with virtual objects. Since the museum's visitors are mobile in physical space, their spatial positions are translated into virtual positions in the electronic space relative to virtual objects. The virtual environment enables the definition of virtual sound sources and the segmentation of the physical space into virtual zones. Invisible to the user there exists a specific category in the broker's domain model for each exhibition object. Depending on the visitor's location, preferences and so forth, combinations of such categories are dynamically selected in order to pre-filter the information items. The next section illustrates how this solution can be adapted to an electronic museum.

Adaptation to a Web Museum

In order to transfer our information brokering solution to the domain of an electronic museum, we present a modus operandi for mapping the physical space into an electronic space. For this reason we follow an overlay model approach for context-aware information systems as introduced by Gross & Specht (2001). This methodology enables us to apply and reuse the algorithm for information filtering described in the previous section within another domain. Thus, only three concepts of the LISTEN application must be adapted or extended: tracking, modeling of the information items, and domain modeling.

As described in the last section, the domain model connects the objects to categories. If the user of the LISTEN system enters the zone connected with a certain domain object, the LISTEN application selects the category associated with this object, in order to filter information items according to the users preferences *and* position. In information brokering terminology, the application builds an information model and uses the broker as triggered by events. In an electronic museum application, the tracking system would deliver the current URI as position and the selected objects (e.g. text links, images, or clicks on maps) as a kind of the visitor's movement. Here, an event is fired when the user visits any Web source that is connected with a category of the broker's domain model.

The information architecture of the LISTEN system is limited to audio presentations. The information broker falls back on a set of audio pieces (e.g. spoken text, music, effects) and returns the best matching item according to the inquired interest model of the user. In the domain of an electronic museum, the curator may probably extend this set of information items by any kind of multimedia presentations, text fragments or links to Web pages. Therefore, the broker's domain model must be extended also to adequately reflect the changed information architecture. New categories have to be added and the meta data description has to be augmented, easily accomplished with the assistance of the information brokering tool. As a result, the system provides additional facility for specification of preferences to the users (e.g. users can turn off streamed video-presentations), and thereby the user model is extended at the same time.

In contrast to the fixed sound installation in the real world LISTEN environment, an electronic museum has to take into account the abilities and features of the visitor's particular environment; i.e. the Web browser. An initially executed piece of software can ask for the user's specific settings and automatically preset the information brokering tool. Based on known environmental properties, the information items are filtered in advance and presented in an appropriate format. The ontology model of the information brokering tool has to be adapted in a suitable way.

The mentioned extensions of the information brokering tool do not affect its capability of filtering information items that fit best the user's context. The overlay model enables the retention of the presented concepts no matter whether the user moves in physical or in electronic space. The next section shows how the presentation of information and exhibits in an electronic museum can further be personalized and adapted to the visitor's needs and behavior. By means of a user model, additional filters can be triggered or refined.

Personalization of a Web Museum

A Web museum environment offers the ability to adapt the order of the presented exhibits and the attached information to the visitor's profile and position within the exhibition. In order to provide a personalized adaptation of the environment according to the visitor's context (i.e. interests, preferences and "motion"), the system has to build up and maintain a user model.

By using sensor data, the user is tracked in the information space in the same way a location tracking system tracks users in physical space. This allows for a variety of new applications and is also an important data resource in user modeling and adaptive systems. Basically speaking, the growing number of data resources about a user allow more valid inferences and much more contextualized interactions between users and adaptive systems.

The enrichment of information items with significant meta information enables the personalization and customization of information offers. By requesting user preferences ,different user profiles can be built up to facilitate information filtering according to the user's needs. Besides information presentation, the system can provide recommendations to the visitor regarding context. These recommended exhibition objects may attract the visitor's attention by emitting an attractor cue (so-called "prompting" of the user). On the basis of meaningful user profiles, several adaptive strategies can be applied to guide the visitor to the Web museum on a specific (strictly predefined or context-aware adapted) tour.

In our approach, the personalization process is divided into four steps: information collection, modeling, controlling and rendering. Each step fulfils a certain role within the user modeling process. The next subsections describe these modules in more detail.

Information Collection

A network of sensors is placed in the environment and connected to variable parameters of the domain. These sensors are used for recognizing changes within the environment, and especially for the perception of the user's interaction with this environment. An observation module receives all incoming events sent by the Web-Server. These event descriptions are pushed into a database. Thus, an event history for every visitor is saved, and an implicit user profile is recorded.

Modeling

By the means of statistical models, the implicit user profile already allows the deduction of valuable information that can be used for standard adaptation activity (e.g. the more time the visitor spends with the art exhibit, the more s/he likes it). This deducted information builds a *behavior model* of the visitor. The behavior model can be treated like explicit representation of implicit user feedback and may be consulted to draw an assumption about the user's interest in a specific object. It is planned to gain more significant information relating to the behavior of the user by implementing different machine learning and data mining algorithms to extract semantically enriched information.

In our user modeling approach for a museum environment, we chose to employ our above mentioned adapted information-brokering tool for modeling *user preferences*. Parts of the domain model of the museum environment are mapped to an ontology model (as shown in Figure 3). The users specify their preferences by simply selecting topics they are interested in from the displayed categories. From these requests different visitor profiles can be built up and stored within the tool. Based on these user selections and on significant meta data describing information items, the personalization process performs a pre-selection and pre–filtering of information offers and customizes according to the users' needs (cf. Pazzani & Billsus, 1997).

The utilization of *stereotypes* is common in adaptive systems. With the aid of stereotypes, the personalization engine defines the observation type of a visitor, and thus, the system is able to adapt the scenery accordingly. Every stereotype causes a different presentation and different strategies (cf. next subsection). These stereotypes are also part of the ontology model of the mentioned information-brokering tool. At the moment, the visitor's classification into one stereotype is done manually by that visitor and cannot be changed automatically during runtime. Additionally, the personalization engine is not able to perform an automated clustering of user profiles and derive new stereotypes from this process.

Controlling

Meaningful user profiles accurately document visitor activity within the museum and can be exploited to adapt the environment in order to support the visitor, to provide personalized information or to invest the exhibition with an even more artistic touch. Therefore, a controlling component is necessary to decide what consequences must occur if certain conditions in the visitor's environment and in the individual user model configuration appear together. Based on these information sources ,the control layer assembles a sequence of commands in order to adjust certain variable properties of the environment. Thus, different sequences of commands lead to different kinds of information presentation. This adjustment of environmental parameters is used for realizing domain independent adaptive or strategic methods, and domain dependant expressive methods. Examples for domain independent adaptation methods on a strategic level are adaptive prompting and adaptive annotation for objects. Expressive methods that seem to be appropriate for the personalization of Web museums are the following:

Adaptation of the presentation

The basis for every kind of adaptation is the presentation of the exhibits with some associated information. Besides the decision about which item is to be shown, the ways presentation can be modified are manifold, and with combinations of these possibilities, a wide range of adaptability is already accomplished.

Adaptation to social context

If visitors are spatially and temporally similar (e.g. two visitors looking at the same exhibit), they may obtain similar information or may be brought together for a chat. Through building such clusters of people, for example, a subsequent discussion about seen objects is possible (cf. Zimmermann, Lorenz & Specht, 2002).

Adaptation to the level of "immersiveness"

Within a museum environment, interest in objects may be expressed by the time a visitor's focus lingers on these objects. The level of interest corresponds to the complexity, the amount, and the style of already received information about one object and is transferred to succeeding objects. If one of these objects complies with the visitor's interests, the presentation style directly steps into the right level of interest, and information items that are classified at the adequate information depth and style are displayed.

Adaptation to movement and perception styles

Several kinds of common behavior can be identified with people "moving" through the environment (e.g. clockwise in real museums). Attractor cues (e.g. sounds, blinking or marked links) emitted from different sources are used to draw the user's attention to certain objects in the environment. Thus, entire tours through the Web museum can be recommended. The selection and dynamic adaptation of tour recommendations can be adjusted to the visitor's stereotypical type of movement and preferred perception style.

Rendering

Rendering means handling the connection back to the domain. This engine translates the assembled sequence of domain-independent commands into domain-dependent commands. The implemented domain-dependent methods directly change variable parameters of the domain (i.e. content of a HTML page) according to the user's behavior. Thus, the decisions taken by the controlling component are to be mapped to real world actions.

Conclusion

Information brokers mediate between information demands and information offers. The information items to be brokered in the LISTEN system are sound pieces the users experience during their movements through everyday environments. The use of an information brokering tool facilitates the management and distribution of a huge amount of information within this domain. Through mapping domain properties to an ontology model, we benefit from a better understanding of the domain's complexity. In addition, the ontology model, in combination with the enrichment of information items with meta data descriptions, enables a personalized presentation of audio information.

In this contribution we presented a concept for reusing our experience, gained during the application of the information brokering tool within audio-augmented environments, in the context of a Web museum. We followed an overlay model approach that builds on and extends the LISTEN methodology in three ways: tracking, information item modeling ,and domain modeling. The information brokering tool supports the adaptation of these three concepts, so that an electronic museum may take advantage of the benefits.

References

Dey, A.K. & G.D.Abowd (2000). Towards a Better Understanding of Context and Context-Awareness. In the 2000 Conference on Human Factors in Computing Systems (CHI 2000): Workshop on The What, Who, Where, When, and How of Context-Awareness. Hague (Netherlands)

Eckel, G. (2001). Immersive Audio-Augmented Environments. *In Proceedings of the 8th Biennial Symposium on Arts and Technology.* Connecticut College, New London, CT.

Goßmann, J. & M. Specht (2001). Location Models for Augmented Environments. In the Workshop Proceedings of *Location Modelling for Ubiquitous Computing*, Ubicomp, 94-99.

Gross, T. & M. Specht (2001). Awareness in Context-Aware Information Systems. In Oberquelle H., R. Oppermann & J. Krause (Eds.) *Mensch & Computer – 1. Fachübergreifende Konferenz*. Bad Honnef (Germany). 173-182.

Jarke, M., R. Klemke, & A. Nick (2001). Broker's Lounge - an Environment for Multi-Dimensional User-Adaptive Knowledge Management. In HICSS-34: *34th Hawaii International Conference on System Sciences*, Maui, Hawaii.

Klemke, R. (2002). Modelling Context in Information Brokering Processes. PhD Thesis, RWTH Aachen (Germany).

Klemke, R. & J. Koenemann. (1999). Supporting Information Brokers with an Organisational Memory. In 5. Deutsche Tagung Wissensbasierte Systeme - Bilanz und Perspektiven, Workshop Wissensmanagement und Organisational Memory (XPS-99). Würzburg (Germany).

Pazzani, M.J. & D. Billsus. (1997). Learning and Revising User Profiles: The Identification of Interesting Web Sites. *Machine Learning* 27. 313-331.

Shilit, B.N., N.I. Adams, & R. Want,. (1994). Context-Aware Computing Applications. In *Proceedings of the Workshop on Mobile Computing Systems and Applications*. IEEE Computer Society, Santa Cruz, CA. 85-90.

Tramberend, H. (1999). Avango: A Distributed Virtual Reality Framework. *IEEE Virtual Reality Conference*. Houston, Texas, USA.

Unnützer, P. (2001), LISTEN im Kunstmuseum Bonn, *KUNSTFORUM International*. Vol. 155. 469-470.

Zimmermann, A., Lorenz, A. & M. Specht (2002). Reasoning From Contexts. In Henze, N. (Ed.) *Personalization for the Mobile World: Workshop Proceedings on Adaptivity and User Modeling in Interactive Systems (ABIS)*. Hannover (Germany). 114-120.

The State of the Art in Museum Handhelds in 2003

Nancy Proctor and Chris Tellis, Antenna Audio, United Kingdom and USA

Abstract

This paper examines the rapidly changing state of museum handhelds from both a technology and content perspective. The paper will also discuss the component parts for a successful museum installation: content, user interface, applications, form factor, positioning and the challenges to integrating these components. Specific examples will be demonstrated from the 2002 Tate Modern installation and the authors' own research.

Keywords: museum handhelds, exhibits, tour guides, Museum Docent™, m-ToGuide

Introduction: The Evolution of the Museum Handheld

In museum technology we are now at a major crossroads with the advent of next generation handhelds (PDAs). As described in previous papers (Tellis and Proctor, 2002; Tellis 2001), the 37 year history of audio guides in museums has seen a slow evolution from reel to reel tape to cassette to digital RAM, MP2 and now MP3 systems. In the first 35 years there were only two major technology shifts, the first being compact cassette in 1980 which greatly reduced the size of the players and the second being the transition from analog to digital systems in 1994.

The transition to digital was by far the most revolutionary as it freed the portable audio experience from the 45-minute time limit and single, forced path to a new experience offering unlimited capacity and free roaming. In the late 90s, with visitors able to choose their own routes and listen as long as they liked, museums began to accept audio guides as an increasingly versatile and desirable visitor accessory. Suddenly, it was possible to provide multiple languages and custom guides for specialized audiences. The opportunities for broad acceptance by diverse museum audiences led to the emergence of audio guides as a standard museum accessory.

Now, at the beginning of 2003, virtually every major and middle-sized art museum in the world has an audio guide for its permanent collection, often in several languages. Acceptance of portable information appliances in museum has also expanded well beyond art museums. We now find ourselves working in 2003 on productions as diverse as guides for The Statue of Liberty, the Spy Museum and the Museum of Sex.

Broad application has also sparked an interest in universal distribution. More than half the audio guides in the world now are bundled into the ticket price and given away to all visitors. Although this is just one of many distribution formats, the combination of random access and universal distribution has emerged as the audio guide format most favoured by museum visitors. Visitors greatly appreciate the ability to control the sequence and pacing of the tours, particularly for the permanent collection. With audio guides normally renting for five and six dollars, the low, universal per-visitor fee, usually around a dollar, incorporated within the general admission fees, gives the impression of a valuable, free educational service.

The development path of handheld computing is now bringing us to remarkable new multi-media players. At the same time, the current generation audio-only MP3 players have become highly sophisticated digital devices with tremendous capabilities. The range of features includes the ability to track and download visitor usage patterns, respond to visitor queries, carry multiple languages, provide linear and random access tours, and synch multi-lingual sound tracks to films and videos. Improvements in compression are dramatically expanding the capacity of the players and challenging wireless for the ability to offer unlimited content. This spring Antenna Audio is introducing a player that will allow two gigabytes of memory, enough for 125 hours of high quality sound. At the other end of the spectrum, RSF has introduced a near disposable low cost wand with minimal features eliminating, for instance, the recharging facility.

Museums are also becoming more sophisticated consumers. Whereas in the past the institution might simply ask for an audio guide, now it will request linear or random access, MP3 players, wands, headset devices, radio systems or PDAs.

A surprisingly energetic debate still flourishes around the subject of headsets vs. wands. Conflicting studies show visitor preference for both options (British Museum Study, 2000; Edinburgh Castle Study, 2001; Museum of Fine Arts, Houston Study 2002; Xerox Parc, 2001). Objections to headsets usually focus on the potential to affect the social aspect of a museum visit. Beverly Serrell (1996), for instance, acknowledges the advantages to learning that comes from simultaneous use of ears and eyes, but notes the isolating effect of headsets. The background for this concern is often based on research conducted by Dr. John Falk and Dr. Lynn Dierking at the Institute for Learning Innovation. Falk and Dierking argue that the educational context is enhanced by the social aspect of the museum visit and that interchange among companions leads to an important form of collaborative learning. They identify the museum as an ideal context for this "free choice learning" wherein visitors absorb ideas and connections at random, based on personal interests and triggered lines of inquiry (Falk and Dierking 2000). Drawing on this research, one might conclude that headsets interfere with this dialogue.

However, recent research conducted at the Xerox Palo Alto Research Center, now known as Parc, has refined this argument in the context of new technologies in museum spaces. In the course of their research into screen-based players, Paul Aoki and Allison Woodruff compared text to headsets to built-in speakers as information sources for handhelds. Their results supported other studies that showed that audio was preferred over text as it allowed the visitor to stay focused on the museum display (Woodruff et al. 2001). Visitors particularly liked sound played over hand held speakers, but this was obviously not practical for indoor museums. The Parc team later tested several kinds of headsets with a group of eight Xerox employees and reported a preference for a single earpiece with headband; sort of a telephone operator rig which combines the hands-free convenience of the stereo headset with a 'free ear' for companions (Grinter et al. 2002).

Clearly a headset-based solution will be the default for screen-based players, as you will not be able to hold them up your ear and also see the screen. But it is still possible to give visitors a choice of headset or wand with their audio tour, as the same players now can be used with excellent sound for either option. The preference of visitors for this choice serves to underline that the museum is a center with many different opportunities for personal learning – some benefit from social interaction, while others require the ability to concentrate and in fact be temporarily isolated from companions. Ultimately, the effectiveness of these educational experiences is based on the appropriateness of the technology used for the stated educational or interpretive goals.

The Centrality of the Visitor Experience

With such a variety of technologies on offer, studies of visitor needs and responses have taken on an even more crucial aspect, not only to inform museums' purchasing decisions today, but also to spur development of visitor interpretation systems for the future. In this spirit, Antenna Audio has introduced the 'TouchPol' touchscreen visitor surveying system into selected museum sites in order to gather visitor response to the audio tour and related issues. In the context of audio-visual handheld systems, CIMI has recognized the rising demand for consistent and effective evaluation tools and is working with Professor Geri Gay and the Cornell HCI Group's concept mapping procedure in the Handscape program "to collect and analyze visitor expectations, museum expectations, potential use scenarios, and the like, as a means of establishing a firm foundation from which future testing may be developed and evaluated." (CIMI 2002)

Similarly, the Multimedia Tour pilot (MMT) conducted at Tate Modern in 2002 took as its primary aim the evaluation of visitor response to the technology and a wide range of approaches to multimedia content design. At the same time, the MMT allowed both Tate Modern and Antenna Audio staff to begin to consider how this technology might fit within the larger interpretation strategy and alongside other interpretation media such as audio tours, wall text, and educational events.

Jane Burton, Curator of Interpretation at Tate Modern, sees the MMT as a sign of times to come for museum interpretation:

> Audio guides have been part of the interpretation and education strategy since Tate Modern opened in May 2000. Having established a reputation for delivering excellent audio tours, Tate Modern would like to remain at the cutting edge of educational technology by helping to shape a new generation of multimedia tours. The advantages of the system we have been testing include the ability to offer content directly in front of the work of art without encroaching on the aesthetics of the gallery space, the potential for creating individualized tours, and for encouraging interactivity (the facility for users to email content home, respond to questions etc). Developing the Tate Modern Multimedia Pilot, in collaboration with Antenna Audio, revealed some pitfalls but predominantly the amazing potential for this new technology, which, in the coming years, will surely take its place alongside more traditional learning tools as a key strand of museum interpretation. (Burton, 2003)

The extensive user feedback gathered during this project at Tate Modern can therefore serve as a lens through which to view both the state of the art in museum handhelds and the near future of museum interpretation, as seen not only by the technologists, but more importantly by the museum's visitors.

The Tate Modern Multimedia Tour Pilot

Sponsored by Bloomberg, Tate Modern's Multimedia Tour Pilot was open to the public from July through September 2002. Developed with Tate Modern by Antenna Audio, the multimedia tour was a 45-minute tour of the *Still Life/Object/ Real Life* galleries, in which visitors could experience audio, video, still images and a variety of interactive applications on handheld iPAQ computers loaned by HP. The content of the multimedia tour was delivered to the visitor through the museum's wireless network, using location-based technologies powered by PanGo Networks. Design support for the project was provided by Tate Modern's Web design team, Nykris.

The pilot had a dual purpose: to test both applications of wireless technology in the gallery, and to assess a wide range of approaches to content design. The wireless network extended over fourteen galleries, defining sixteen unique content zones by means of just seven access points. In each gallery, a different approach to content design was applied to the presentation of a selected object. After their tour, visitors were asked to evaluate the different designs and provide feedback on the overall impact of their experience of the technology.

The pilot was specifically designed to define the limits of the current generation of wireless technology, and therefore involved an ambitious combination of location-based and interactive applications, namely:

- **Interactive survey and response:** the system asked and recorded visitors' opinions of a Nikki de Saint Phalle painting, both before and after they saw a video of its production;
- **Creative play:** visitors could mix their own soundtracks to accompany their viewing of visual collages by Edoardo Paolozzi;
- **Location-specific content delivery:** content was delivered to visitors according to their location in the gallery.
- **Visitor tracking:** the Proximity Platform™ tracked visitors' locations in the galleries, relaying them to a digital map at the distribution desk. This system also alerted staff if a device failed or the visitor otherwise appeared to be in need of technical assistance.
- **Visitor paging:** staff could page visitors *ad hoc* through the PDAs, as well as send out pre-programmed, timed alerts, e.g. regarding the start of a video program in the gallery.

- **Visitor profiling:** the wireless system tracked visitors' use of the content with a record of which exhibits they visited.

- **Visitor e-mail facility:** visitors were able to e-mail themselves further information on objects and artists on the tour in order to follow up on artists and artworks of interest through the Tate Web site.

Visitor Feedback

A wealth of visitor responses was gathered from the pilot: 852 visitors took the pilot tour and completed evaluation forms recording their experiences. These responses were analyzed in a comprehensive database. In addition, the software system used in the trial logged all uses of the MMT and provided a statistical picture of how the tour was utilized, which rooms were visited, and how the visitor e-mail system was used. Not least, Antenna Audio staff demonstrated the tour to a wide range of visitors and specialists, including curators, educators, museum directors, artists, technologists, and wireless solution integrators. In conversation with these experts in their fields, we gained valuable and detailed feedback on the MMT technology, the experience, and ways of improving both.

Although this technically innovative pilot often pushed the technology to its limits and beyond, visitors were enthusiastic about both the service and the tour. Visitors generally see this technology as an exciting and inevitable part of the future landscape in museums. The British Academy of Film and Television Arts agreed that the multimedia tour enhanced the visitor experience at Tate Modern, remarking in its award to Tate Modern and Antenna Audio for Technical Innovation that:

> Genuinely groundbreaking, this was an exciting demonstration of how new technology can be used to enhance museum and gallery visits. Using a hand-held wireless device that knows just where you are on the tour, this offers a stimulating array of material to add to, but not confuse, the experience of a gallery visit. Commendably, Tate Modern is working with day-to-day feedback from visitors to develop a system that complements an already stunning physical learning space. (BAFTA 2002)

In this respect, the trial was a resounding success in demonstrating how such cutting-edge technology can be implemented today in a museum's interpretation program. It achieved its goals of mapping both the strengths and weaknesses of wireless handheld solutions; however, the true judges of the success or failure of this new technology will be the museum visitors. As a result, the visitor responses gathered during the trial are our most valuable roadmap towards development of handheld technologies for museums in 2003 and beyond.

Visitor Demographics

A primary tool in the collection of visitor response to the MMT was a questionnaire that was completed by all visitors taking the tour: 852 questionnaires were completed by visitors from around the world.

- The largest group of visitors fell in the 26-40 age bracket, with 26% of visitors aged 18-25, 24% aged 41-60, 9% aged 10-17 and 4% over 61.

- 42% of visitors were female, and 58% male;

- 56% of visitors were British;

- 18% were North American;

- 17% were from Continental Europe.

- 8% were from the rest of the world (1% did not respond to this question)

The comparison of this sample with the overall demographics of visitors to Tate Modern indicates that audio-visual tours will be particularly popular with North Americans:

- The MMT attracted more than double the percentage of North American visitors at Tate Modern in this period (18% of MMT users were North American, compared to an overall visitor ratio of only 7%);

- The MMT attracted fewer UK visitors than the average in this period (56% of MMT users were

British, compared to an overall of 70% UK visitors to Tate Modern);

- The MMT attracted a similar percentage of European and 'Rest of World' visitors to the Tate Modern standard;

- The MMT attracted more male visitors than the average at Tate Modern: 58% compared to a usual audience of 41% men at Tate Modern;

- The MMT attracted fewer visitors aged over 60 than the Tate Modern average, but 3-5% more visitors in the under-60 age bracket.

Visitor Satisfaction

The average amount of time visitors spent taking the tour was 55 minutes.

Over 70% of visitors spent longer in the gallery, and a similar percentage said that the MMT had improved their visit to Tate Modern.

Approximately 61% of the visitors gave their email addresses in order to be able to use the e-mail feature. The overall content rating the tour received was nearly 7 marks out of 10. The clear conclusion from user response to the tour was that visitors enjoy the tour more when the content is tailored to their interests and ability level. As a 'sampler' of content design approaches, the MMT was geared for a generalist audience with interest in the art on display, but no prior knowledge. Unsurprisingly, then, we found that visitors with extensive art backgrounds were less satisfied by the content on the tour, and would have preferred more research-oriented content and facilities.

Main Difficulties

In general, older visitors found the technology more difficult to use than younger visitors. Overall 55% of visitors found the MMT easy to use, while 45% found it difficult.

Regardless of age, nearly all visitors experienced technical difficulties in using the new technologies of the MMT. The most requested improvements for the system were making the content easier to access (i.e. by improving the user interface), and making the PDA more reliable.

Technical Findings

The technical performance and ease-of-use of a handheld tour are in fact determined by seven different elements that must work together in order to provide a positive user experience. These are:

The PDA, including:

- **Processing Speed:** the current generation of PDAs have very slow processors, and are roughly comparable to 6-10 year-old desktop computers. They are the 'bottleneck' in the wireless system's ability to delivery media-rich, processor-demanding content to visitors;

- **Operating System Stability:** PocketPC was not originally intended for rich multimedia content and so is not yet optimized for this use; however, there is huge pressure from the PDA manufacturers on Microsoft to make Internet Explorer 5 available on the Pocket PC. A major rewrite and upgrade is expected within the next 6 months (.NET). In addition, alternative operating systems are available, including Pocket Linux, Palm OS 5 and C++ content delivery engines.

The Wireless Network

The 802.11b protocol used at Tate Modern provides an embarrassment of bandwidth - effectively 5mB per second data transfer rate (up to 11mB per second in lab conditions). This is more than twice the speed of a T1 Internet connection. However, the processor in the PDA was not always up to handling the quantity and complexity of files being delivered to it. Delays, therefore, in the playing of content, etc., were due less to network issues than to processor weaknesses.

'Client Side' Software

'Client side' software is stored locally on the PDA, and includes Internet Explorer for PocketPC 2002 and Macromedia Flash v. 5, the file format in which the tour's multimedia content was programmed.

'Server Side' Software

'Server side' software is stored on the network server, including the content delivery engine (PanGo's Museum Docent™) and its applications that allowed us to create content in industry-standard formats such as HTML and Flash and deliver it to the user. The Museum Docent™ software and its applications proved to be very stable; indeed, the software did not crash during the entire pilot period.

Special features of the Museum Docent™ software include:

- **Full Screen Content Mode:** so that it was not necessary to sacrifice any screen real estate to the Internet Explorer browser window;
- **Paging Service:** allowing the distribution desk to contact visitors using the tour through the PDA;
- **Alerts Service:** automatically prompting the visitors before each start of the Jan Svankmajer video sequence;
- **e-Phile Service:** allowing visitors to bookmark information of interest and e-mail it to themselves.

The Location-Based Service (LBS)

The LBS provides location-sensitive content delivery and location-sensitive interactive maps of the gallery. PanGo's Proximity Platform™ calculates the location of the user by taking signal strength readings from a number of access points in the user's vicinity and effectively 'triangulating' the user's position. In the MMT pilot, we were able to define 16 distinct content zones by using only 7 access points - a significant cost saving over earlier location-based systems in terms of network installation and maintenance. Limitations of the Proximity Platform™ include:

- Minimum location granularity of 1 meter in ideal conditions; in practice, 3 meters;
- Ambiguity in 'border spaces' and overlapping content zone areas;
- 2-3 second latency for accurate location identification.

Design

Design issues include three aspects of the tour:

- Tour Experience Design;
- Tour Interface Design;
- Tour Content Design.

Each design element has not just a visual impact on the user, but also a logical impact: just as each message is structured according to an interpretive or pedagogical logic, so the navigation icons and their arrangement on the screen reflect basic assumptions about the use and the organizing principles of the tour experience. Together the interface and content should support the desired Tour Experience. It is important when evaluating feedback from users to distinguish between these two elements, and also to keep in mind that the tour was not conceived as a united 'Experience'. Instead, the purpose of the pilot was to test a variety of approaches to the content design, as well as the technology and visitor feedback to it.

Human-Machine Interface

The form factor and weight of the PDA, as well as the screen and battery life, also impacted the way the visitors interacted with the MMT. In particular, the following aspects presented challenges to the operations of the pilot project:

- **Battery life and charging:** The iPAQ's screen drains enormous power from the PDA. Although HP provided very helpful guidance in setting up the PDAs for maximum power saving, we still had to leave the screen on constant, high brightness in order to provide the best possible user experience. As a result, the battery power and therefore performance of the iPAQ deteriorated after approximately 1.5 hours of use. The iPAQ would then need to be recharged for approximately one hour in order to reach full battery power again. This severely limited the number of uses each iPAQ could have in a day.

Smaller, lighter devices with integrated wireless cards and greater battery life significantly improve the user experience of the PDA. The batteries used in the

3900 series iPAQs and the new 5000 series provide both longer life and better power management. HP's new X-Scale-based 5000 series iPAQs with integrated wireless include a removable lithium polymer battery, so it's possible to keep spare batteries charged up in order to minimize down time during recharging. These iPAQs also include the better quality screen that is already in the 3900 series iPAQs, and the stronger X-Scale processor. Toshiba, Acer, and Symbol also offer PDAs with integrated wireless cards which improve not just the battery life but also the ergonomics of the device. Antenna Audio is currently developing mass charging facilities that will be compatible with industry-leading PDAs.

- **Weight:** Many users found the iPAQ too heavy to carry for long periods. Although the large straps distributed the weight more evenly across users' necks, the width made them uncomfortable with many people's clothing. Small children were not able to use the devices for very long because of their weight.

- **Fragility:** In order to protect the PDA, it was given to the visitor in a rugged case which increased the weight of the PDA. However, we found that this level of ruggedization was sufficient to protect the device. In three months of use, we had no breakage, though several devices were dropped.

- **Screen:** In general, users are happy with the size of the standard PDA screen, as long as the content is designed well for it. In particular, content design needs to take into account the screen's 'high contrast' structure, designed to make the screen legible out doors, which can render extreme darks and light colors illegible. Resolution also helps the legibility of the screen enormously, and we saw a marked improvement in the screens used in the 3900 series.

Content Findings and Recommendations

In addition to testing the technology and visitors' responses to it, the primary aim of the MMT pilot was to test a variety of approaches to content design. The content proved to be the primary draw of the MMT, and indeed it will be the quality of the content that ultimately determines the success or failure of the tour experience. In many respects each new object demands a new approach, so accumulating a design guide of what works and what doesn't is likely to take as long if not longer to develop than the technology.

The findings and recommendations made below regarding content design are based on Antenna staff experiences, the questionnaires, and feedback from Susie Fisher, the independent evaluator hired by Tate Modern.

What Worked

Interestingly, users did not seem to find multi-tasking and multi-tracking of different media (e.g. looking between screen and artwork) to be a problem as long as the message was well designed and the PDA was functioning properly. The multimedia tour clearly had the effect of making the visitors look longer at an object than they would have otherwise, even though the screen was also commanding attention. As Susie Fisher reported, "Visitors can multi track with great ease, even when the input tracks (audio, screen, painting) are not synchronized with one another." (Fisher, 2002, Chart 34)

In this regard, "audio acts like a friend", and indeed more use could be made of the audio to direct the users' eye movements between the object, and the screen, and to navigate through the gallery space.

In both the questionnaires and Focus Groups, visitors' favorite stops on the tour featured the following design approaches:

- **Audio-Visual Coherence:** A strong logical link between the audio and the visual worked well in messages where an audio description of abstract paintings by Braque and Auerbach was accompanied by visual 'zooming' into the details being discussed; however, it is important to note that at least in the case of the Braque message, the 'audio focusing' of the message that explained the painting's various elements could have been accomplished with audio alone. The Auerbach painting, on the other hand, is more abstract and therefore really benefited from the explanatory visual details. Susie Fisher's advice here is apt: "Use the visual to do things which can only be done visually" (Fisher, 2002, Chart 32).

- **Interactives:** Interactive messages, in which visitors had a chance to register their opinions about a *Shooting Picture* (by Niki de Saint Phalle) and create their own soundtracks to listen to while viewing collages by Edoardo Paolozzi, were very popular. Visitors asked that the interactivity be enhanced further, for example, by showing them how other visitors responded to Saint Phalle's work, and by allowing them to save or send their Bunk soundtrack combinations.

- **Audio:** Visitors responded enthusiastically to messages that included interviews with artists, sitters, and related experts, as the speakers were relevant and interesting. A porno-movie-style soundtrack by the Chapman Brothers entertained with its shock effect, but it also alienated other visitors. It became clear that good audio navigational instructions are also a core skill, required by both the MMT and the traditional audio tour, since moving visitors around the galleries safely and effectively is fundamental to the success of any tour.

- **Video:** Perhaps surprisingly, video was not the focus of praise for the message that showed artist Niki de Saint Phalle creating a *Shooting Painting*. Instead, visitors most commented on how much they enjoyed the process of learning in this message, where they had the opportunity to give an opinion of the work both before seeing the video, and afterwards. Using the screen to explain the process of making a work was considered by several visitors to be a good use of the screen, but the video was also considered a potential distraction;

- **Intuitive, Interactive Interfaces:** A message about Damien Hirst's *Pharmacy* used a 360 degree panorama of the installation as the interface for audio messages about the work, much in the style used by the Parc researchers at the Filioli Mansion. Although this room was de-installed early on in the project and therefore was not in the focus group analysis, it was very popular among visitors for the ease of use of its interface as well as for the interviews with the artist and a pharmacist which it included.

What Didn't Work

- Long Messages: Visitors seemed to tire more easily from the interactive tour than from more passive audio tour experiences. In addition, the screen seems to make them expect more to happen, more quickly, than in an audio tour. It is likely, therefore, that multimedia messages need to be shorter than audio tour messages, but to offer the visitor more second level message options. The PDA screen, which can display several second level options, makes this possible, in contrast to audio tours where it is more difficult to offer multiple message 'layers' because of the difficulty of remembering their number combinations. However, as in audio tours, second level options can be a useful way of dispersing audiences by directing them to related works, e.g., so that they don't spend a long time bunched in front of a few selected works in each gallery;

- Blank Screens: We intentionally made the screen go relatively blank at selected points along the tour, in order to focus visitor attention on the art object. Visitor feedback indicated that we were probably too conservative in this: they found the blank screen confusing, and wondered if something were wrong with the PDA. The screen demands that at least a still image of the object being discussed appear on the screen at all times, as much to identify the object of discussion and reassure the visitor that the PDA is working as anything. Regarding our fears that visitors would spend their time in the galleries 'with their eyes glued to a small television screen', one visitor comment seemed to sum up the lack of danger: "there is no way the small digital image of the object could be more compelling than looking at the real thing." Here, as ever, good content design is key to the tour, producing a beneficial rather than distracting visitor experience;

- Text: Text had a mixed response: some - particularly more 'art experienced' visitors, liked having wall labels in the palm of their hand, while others wanted more exciting content;

- Help Menu: The lack of a help menu was felt, as visitors needed to be able to refer to a key to the navigation icons throughout their tour, to remind them of the functions and options available to them.

And Moreover

- Visitors wanted MORE of everything: more objects on the tour, and more information about each.

- Just as in audio tours, the multimedia tour can take attention away from other objects in the gallery, items which are not on the tour. Therefore careful tour design is essential.

User-Interface

Interface buttons and Functions

The MMT used the following interface buttons:

- Map
- Information
- What's Here
- E-Phile
- Rewind
- Pause/More

The icons and functions of these buttons were often quite unclear to users, and there were several functions that users missed and requested. This response pointed to the fact that this new technology requires new functions that are not yet known, let alone intuitive, to the general public. In response to this feedback from the Tate pilot, Antenna Audio is redesigning the interface and multimedia tour functionality, and will be working with Tate Modern to refine these further.

Content Menus

In order not to interrupt or control the visitors' experience of the tour, the MMT was designed to leave the launching of content entirely in the hands of the visitors; although the content could have been triggered by the location of the visitor to play automatically, the visitor had to press the 'What's Here' button or choose a room from the map in order to trigger its content to play.

We found, however, that there was demand for an automatically-generated menu of content available in each gallery, to both respond to the visitor's movements and prompt the visitor to explore new content. This demand can be satisfied through the Proximity Platform's™ 'aggressive' location sensing mode. When a message stops playing on the device, and the user is not receiving any content, the Proximity Platform™ will go into 'aggressive' mode, sending a Local Area Menu of information relevant to the user's immediate area to the PDA's screen. This menu will not interrupt any content the user is currently playing; it will only appear if the user is not currently playing content. Alternatively, users can specifically request the content menu for their local area at any point in the tour.

Customer Service Findings and Recommendations

Distribution Desk and Management Console

At the distribution desk's computer terminal, the Museum Docent™ displayed a map that tracked the location of users in the gallery as they took the tour. This and the 'checking in/out' system, which allowed staff to associate a visitor's name and email address with the PDA, and hence with the location on the gallery-tracking map, were particularly useful for security and tour management. Similarly, the server-side usage logs, which tracked where visitors went, which messages they accessed, etc., proved particularly helpful in evaluating the success of the pilot.

The Museum Docent™ also includes a client-side log, which sits on each PDA and tracks any problems that occur and why. This performance log proved essential to monitoring PDA performance.

Distribution and End-User Training

The MMT was emphatically represented to the public as a pilot; by being given the tour for free, and being asked to fill in evaluation forms, the public was reminded that the aim of the installation was to gather information in order to improve and develop the service further.

The best way to distribute the PDAs and teach visitors how to use them was a constant work in process. Even the distribution desk underwent various permutations, as various storage and presentation facilities became available.

The navigation button icons were difficult for visitors to understand and remember, and posed a particular challenge to instructing visitors in the use of the device. It could take upwards of 5 minutes to teach each visitor how to use the tour.

Visitors' credit cards or other major documents were held as security for the PDA. With only 15 PDAs, keeping track of these did not pose any difficulties for distribution staff.

Security

Despite the lack of equipment security at the distribution desk, there were no equipment thefts during the trial. Nor were there any breakages.

The location-sensitive map on the distribution desk proved useful in tracking devices. On one occasion, a couple clearly left the tour area and 'dropped off' the tracking map on the distribution desk. However, because the wireless network coverage extended beyond the tour area, staff were able to send a page to the couple, asking them to return to the distribution desk. They did so and returned the devices, explaining that they had simply gone to the café and had forgotten to bring the PDAs back first.

A Glimpse of the Future for Handhelds and Museums: The m-ToGuide Project

The m-ToGuide (mobile tourist guide) project is funded by the European Union as part of the IST (Information Society Technology) programs, 5th Framework. In 2001, the IST commission made a total of +25 million available for 5 projects to be conducted towards the development of next generation telephone technologies and businesses in Europe. The aim of these projects is to develop technologies and build business for next generation telecoms, by creating consumer demand for new services and hence jobs in the sector.

Motorola's research and development division, Motorola GTSS, headquartered in Israel, assembled the m-ToGuide group, a consortium of 17 companies including Antenna Audio, to respond to this call. The team was awarded a +5.5 million grant in April 2002. The m-ToGuide service will integrate tourist content and services from a variety of providers, delivering these to the end-users' handheld terminals (PDA) via 2.5 and 3G telephone networks.

Currently, 3G networks are not truly available for the m-ToGuide service, so initial trials will be conducted using a combination of locally-stored content and the slower GPRS or 2.5G network infrastructure to transfer low-bandwidth data to the terminal. Positioning will be provided by GPS, but the system is designed to receive location information from a variety of other inputs as well, including cell id, the system used by the current generation of phones (GSM). Cell id is the least precise positioning system available, locating the user only within approximately 500 sq meters, while 'assisted GPS', using GPS plus positioning software, can achieve up to approximately 3 meters accuracy.

The m-ToGuide service is intended to replace and enhance the navigation and information services currently available to tourists, via a handheld device, a kind of PDA, connected to high-speed telephone networks (GPRS or 2.5G and UMTS or 3G). Equipped with location-sensitivity (GPS and cell id), the m-ToGuide terminal will provide interactive navigation services and personalized information, according to criteria established by the users, on places of interest, monuments, stores, restaurants, hotels, museums, opening hours, etc. It is also being developed to include m-commerce facilities, allowing the users to buy tickets and make bookings. Trials of the service will be carried out in Madrid, Sienna and London in summer 2003.

Antenna Audio is responsible for developing audio visual and textual messages for the main points of interest of each city in the project. This 'highlights tour' of each city may be supplemented by other tours, for example an audio-visual tour of a museum, current exhibition, or other visual attraction. The tourist can then use a single handheld device to tour both interior and exterior spaces, receiving both media-rich interpretation of features and attractions as well as practical, navigational support from maps and visitor information. A primary aim of Antenna Audio's involvement in this project is to use the technology to develop new audiences and means of reaching those audiences, as well as new revenues and commercial opportunities, for our museum partners.

In addition to providing the London tour content, Antenna Audio led the research that defined users' needs and requirements for this new city tour technology. In summer 2003, Antenna Audio will put its learning from this and the Tate Modern pilot into action as it staffs and runs the distribution operations and user support services for the London field trial. As in the Tate Modern MMT, user surveying and evaluation will be a main activity and goal for the m-ToGuide project.

Conclusion

The development of next-generation museum handhelds is being informed by the experience gained from millions of audio guide users over the past three decades. The fundamental benchmarks for successful mobile learning experiences are now well known, allowing us to concentrate on new features and opportunities.

Wireless interactive systems offer important tools and unique opportunities for the development of in-gallery interpretation and education programs, and the extension of these to a cultural experience of the wider city that links many museums and visitor attractions. The visitor response to the Tate Modern Multimedia Tour Pilot has given an unequivocal green light to future development of these handheld solutions within the museum. We now look forward to the results of the m-ToGuide project as cultural tourists guide the development of mobile, interactive technologies beyond the museum's wall as well.

References

Antenna Audio (2000) British Museum Study 6/00 showed a 79% preference for headsets.

Antenna Audio (2001) Edinburgh Castle Visitor Survey 2/01 showed a 73% preference for headsets.

Antenna Audio (2002) "TouchPol" survey at the Museum of Fine Arts, Houston, 11/02 showed a 67% preference for headsets.

BAFTA (2002) British Academy of Film and Television Arts Interactive Entertainment Winners 2002. Technical Innovation. Consulted January 15, 2002. http://www.bafta.org/5_ie/winners2002.htm#ti

Burton, Jane, Curator of Interpretation, Tate Modern, Interview, 2003

CIMI (2002) HANDSCAPE - Handheld Access to the Museum Landscape. Last updated October 14, 2002. Consulted January 15, 2002. http://www.cimi.org/wg/handscape/handscape_long_desc_1201.html

Falk, Dr. John and Dr. Lynn Dierking *Learning From Museums*, Alta Mira Press, Walnut Creek, CA 20000.

Fisher, Susie (2002) Tate Modern Multimedia Tour Evaluation, unpublished findings, 2002.

Grinter, R. E. and A. Woodruff (2002). *Ears and Hair: What Headsets Will People Wear?* Extended Abstracts, ACM SIGCHI Conf. on Human Factors in Computing Systems, Minneapolis, MN, Apr. 2002, 680-681.

Serrell, Beverly (1996) *Exhibit Labels*, Alta Mira Press, Walnut Creek, CA 1996.

Tate (2002) Unpublished survey results. Research period 25 July to 02 Sept 2002. Statistics based on a sample size of 445 visitors out of total visitor figures of 677,307 during the research period.

Tellis, Chris (2001) Panel Presentation, Museum Computer Network Conference, Cleveland, November 2001.

Tellis, Chris and Nancy Proctor (2002) Workshop: *Handhelds in Museums*, Museums and the Web, Boston, 2002.

Woodruff, A; M. H. Szymanski, P. M. Aoki and A. Hurst (2001). *The Conversational Role of Electronic Guide Books*. Proc. Int'l Conf. on Ubiquitous Computing, Atlanta, GA, Sep. 2001, 187-208. (© Springer-Verlag.)

Designing Multi-Channel Web Frameworks for Cultural Tourism Applications: the MUSE Case Study

Franca Garzotto, Politecnico di Milano; Tullio Salmon Cinotti, Università di Bologna; and Massimiliano Pigozzi Casalecchio di Reno, Bologna, Italy

Abstract

A framework for the design of multi-channel (MC) applications in the cultural tourism domain is presented. Several heterogeneous interface devices are supported including location-sensitive mobile units, on-site stationary devices, and personalized CDs that extend the on-site experience beyond the visit time thanks to personal memories gathered during the visit. The design framework is multi-layer in nature: it takes care of application design both at conceptual and implementation level. At conceptual level it supports content, navigation, interaction and presentation design. At implementation level it includes an interface-independent execution engine as well as a set of tools mapping the design into a formal interface description that specifies the run-time rules-of-behavior to the execution engine. The proposed framework is going to be demonstrated in outstanding museums and archaeological sites in Italy.

Keywords: multi-channel applications design, pervasive computing, cultural tourism, context-awareness, zero-interface.

Introduction

The Information Society expects the establishment of pervasive services (Abowd 1999, Weiser 1993) providing friendly and effective *education* to everybody, in every place, at any time, improving the *mutual understanding* between people of different countries, and impacting our approach to *global culture*. Cultural tourism offers a precious opportunity to verify the potential and the impact of this emerging technological scenario. The approach of today's tourists to culture has significantly changed with respect to the past. The primary goal of the travellers in the XVII century was "seeing" and "wondering". For travelers of the third millennium, "learning" is usually one of the main objectives.

Information technology offers new challenges for the creation of an education environment that is based on cultural heritage and addresses the *need for knowledge* of cultural tourists (Oppermann 1999, Bocchi 1999, Guidazzoli 2000). In particular, the technology of *multichannel Web applications* (MC applications for short) promises a successful approach to these goals. An MC application delivers its services on several stationary or mobile devices or "channels"; such as desktop PCs, notebooks, mobile phones, PDAs, Web TVs and dedicated appliances. An MC application relies upon a centralized application environment which manages a shared information base coupled to the so-called "business logic" (Salmon Cinotti 2001). On each channel, it delivers services which are *customized* to the characteristics of the delivery device and of the user context (Ozen 2001, Perkowitz 1999). As such, MC applications offer the possibility of making cultural services available on-demand wherever fruitful and convenient for a tourist, in a way appropriate to the physical and logical situation of use.

The state-of-the-art technology provides a number of hardware and software solutions for implementing MC applications. But few methodological tools exist to help developers master the complexity of *designing* MC applications.

The MUSE project – funded by the Italian government under the leadership of Ducati Sistemi S.p.A., in cooperation with Cineca, Politecnico di Milano, Sinet and the University of Bologna - addresses the above issue in a specific domain: *cultural tourism*. One of the MUSE goals is to develop a "framework" for MC cultural tourism applications. According to software engineering, a framework "...provides a reusable solution for a class of software applications that share a common set of requirements in a given domain" (Johnson 1997). A framework can be regarded as an application "skeleton", which captures the essential features of a class of applications and can be "easily" instantiated to produce (some aspects of) a specific application in the class (Fayad et al. 1997, Schwabe et al. 1999, Garzotto et al. 1999).

The MUSE framework has different levels of abstraction. At the highest level, the MUSE framework supports the *conceptual design* of MC applications for cultural tourism. The MUSE *conceptual framework* defines the user experience for the four MUSE channels. It provides the general classes of information and navigation structures that may be shared among different applications in the domain of cultural tourism, and it defines the general navigation and interaction modes by which these contents can be used on each different channel.

At the lowest level, the MUSE framework focuses on the *technological infrastructure*, by defining the general characteristics and integration features of the MUSE delivery channels.

At an intermediate level, the MUSE framework considers the *software architecture*. It defines the middleware components that support the communication between the different MUSE channels and the shared information base.

The MUSE framework, at all of the above mentioned levels, has been tested and used for the (on going) development of three MC applications for cultural tourism: one concerning the archaeological site of Pompei, the second concerning the Certosa di San Martino (a large museum located in an XVII century monastery in Naples), and the third concerning the Florentine "Museo di Storia Della Scienza" (Museum of Science History).

The development of the MUSE MC application in Pompeii is carried out in co-operation with a team of archaeologists currently working in the project "Pompei - Insula del Centenario (IX,8)". This project started in 1999 as a co-operation between the "Dipartimento di Archeologia dell'Università di Bologna" and the "Soprintendenza Archeologica di Pompei". Its primary goal is the study and restoration of a large Roman "domus" (house), called "*Domus del Centenario*", located in the "*Insula*" (Quarter) IX. The MUSE MC application focuses on the *Domus del Centenario*, on its insula, on two routes, both including the magnificent forum, which connectss the domus to the site main entrance and one of the beautiful houses located along these roads. The application contents include maps, audio-video material, images with their descriptions, 3D static and interactive models, and virtual reconstructions of large sections of *la Casa del Centenario* as it is today and as it was in 79 AD when Pompei was destroyed.

La Certosa di San Martino in Naples is a large XIV century charterhouse located on top of the Vomero Hill. It was deeply modified in the XVII century and it became "the Museo of the City of Naples" in 1867, after the abolition of all Roman Church properties ordered by the government of the United Italian Nation. The museum is partitioned into many sections, mostly located in the rooms where the Carthusian monks used to live. The monastery includes cloisters, gardens, and a church. It offers unforgettable views of the gulf, the city, and Mount Vesuvius, and is one of the most beautiful monuments in Southern Italy. The Muse MC application covers a museum section called "Immagini e Memorie della Città" ("Images and Memory of the Town") as well as the areas connecting this section to the main entrance. The section was chosen by the museum authority because of its magnificent panorama and of its link to the city and its history. The MUSE MC application contents include text, images, and audio-video clips about paintings, historical cartography, sculptures, ceramics and furniture.

The "Museo di Storia della Scienza" in Florence is the most important History of Science Museum in Italy. It includes an outstanding collection of instruments paving the path of science evolution in many areas, including astronomy, sky physics, mechanics, optics, magnetism and electricity. The Muse MC application aims to offer an engaging education experience to communicate the meaning and the history of the museum exhibits, and in its first version will cover two halls devoted to Galileo Galilei, focusing on Galileo's instruments, his work, and his scientific achievements.

The rest of this paper will discuss the characteristics of the MUSE channels and will present the main feature of the MUSE framework.

Muse Channels: General Features

The MUSE MC framework considers four different channels: the conventional Web channel, the on-site mobile channel, the on-site stationary channel, and the memories channel. The Conventional Web Channel

This channel is a conventional internet connected desktop or laptop. The user experience of an MC application on this channel is typical of any conventional cultural Web site - navigation based exploration and fruition of cultural contents- suitable for a user who is planning a trip or just wants to learn more about a cultural subject either at home or in a museum entrance or reading room.

The On-Site Mobile Channel

The on-site mobile channel is a context aware hand-held device intended to be used during the visit to a cultural site, e.g., a museum or an archaeological settlement. The mobile device, called Whyre, is specifically designed to act as a personal multimedia interactive guide for the visitor. Whyre, shown on fig. 1, is worn with a neck strap like a camera's. It is connected to a local server by air and it is equipped with a high brightness 6.4" display. "Whyre" enabling technology is an Intel Architecture based platform with dynamic voltage and frequency scaling developed within the Muse project with the kind support of Intel Labs.

Figure 1: Using Whyre, the MUSE on-site mobile channel

Due to the screen and performance limitations of the device, communication effectiveness is based on words, images, and animation. Words are mainly conveyed by audio and a limited amount of text (captions, directions, dates, or on-demand textual copy of the audio clip).

Whyre can work in two main "modes". In *Web mode*, the information base can be accessed as any standard Web site; e.g., by a visitor quietly sitting down somewhere on-site, using the conventional Web navigation paradigm. In the *context aware mode*, the device can detect the approximate user position and field-of-view in the physical space. In this mode, the application can be dynamically customized to present the multimedia content and interaction capabilities which are more appropriate to the current user position.

Context-aware customization is triggered by orientation- and location-dependent events which are identified by properly fused data originating from many on-board sensors, including a GPS, a digital compass, gyroscopes, accelerometers, a camera and a WLAN based tracker. Context-aware customization supports *serendipity*: it makes the visit full of surprises, stimulates resource discovery, stirs up questions, and provides answers at incremental depth levels. In addition, context-awareness supports the adoption of the "virtual guide metaphor": it enables the application to guide the user along site tours designed by the site curators, keeping track of where the users are in the tour, warning them in case of deviation, and suggesting how to return to the tour path.

The On-Site Stationary Channel

The content on Whyre is enhanced by the *on-site stationary channel*, a stationary high performance graphic station connected to a large high-resolution display located at specific points of a cultural site. The graphic station is radio-controlled by the same device used in the mobile channel, and the connection to the graphic station is automatically established when the user approaches the station's large display.

The high performance of this channel and its screen size make this channel appropriate for the delivery of cultural content where space rendering, high qual-

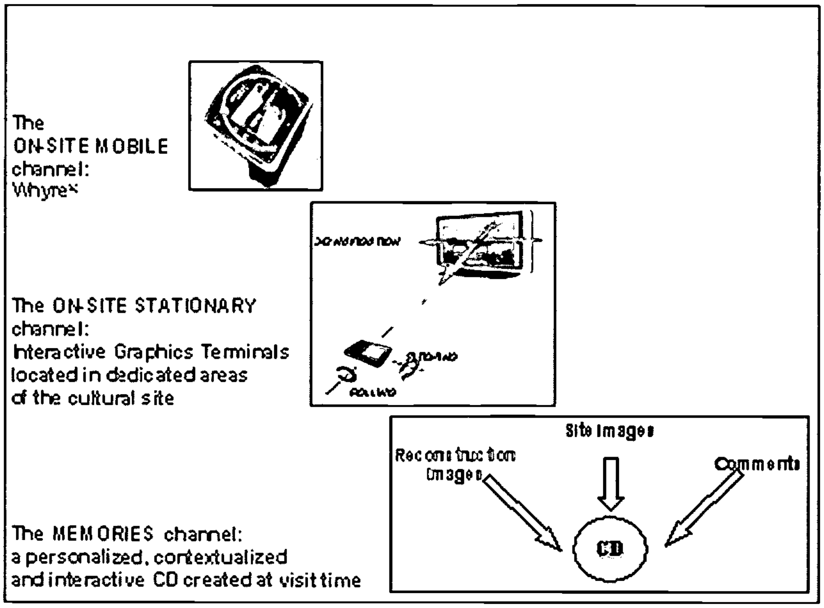

Figure 2: Current MUSE Channels

ity interactive multimedia, and multiple simultaneous visual presentations are key factors.

The user can access the on-site stationary channel in any place where the high-resolution display is located, and can view on the large screen channel multimedia contents that are related to that place: high quality images, 3D reconstructions and simultaneous rendering of multiple synchronised 3D models (comparing, for example, the virtual reconstruction of a building in its "original reality" and the model of its present shape). Users can interact with multimedia elements by using the Whyre control elements or by simply rolling and pitching the hand-held device: this may be appropriate, for example, to move within the 3D space and explore the virtual reconstructions. This channel supports an immersive and exciting cultural experience, and takes the visitor to an easy and engaging learning experience (Bocchi 1999, Guidazzoli 2000, Strohotte 1998). In addition, by switching from the on-site mobile channel to the on-site stationary channel, visitors can turn their individual activityinto a potential "group experience".

The Memories Channel

The memories channel is a multimedia CD-ROM which users can fill in during the visit and bring back home. During the on-site visit, the users can select ("mark") any cultural material presented by the on-site mobile device or by the on-site stationary device, and they can take digital pictures using Whyre to recall the emotions of the visit. As these contents are selected or created, they are sequentially stored in a Memory Album and interlinked. The result is a very simple hypertext which provides a preview of all selected elements, a direct bi-directional link from the preview to each element and from each to the next and to the previous one. At any time, either during or at the end of the visit, the users can preview the album, add comments, delete some items, modify their order. The final album can be saved on the CD, which represents the "memory" of the on-site visit.

So far, the work of the MUSE project has focused on the three channels depicted in figure 2, since they are the most innovative and technically challenging. The conventional Web channel will be included in the framework in the future.

The Muse Application Design Framework

The goal of the MUSE application design framework is to define a generalized conceptual model which captures the most relevant features of MC applications for cultural tourism and can be reused for designing, in principle, any specific MC application in this domain with a minimum amount of refinement and specialization.

The work which led to the framework definition started from the analysis of the application requirements of the three MUSE case studies: Pompei, Certosa di San Martino, and "Museo di Storia della Scienza di Firenze". Each individual case has some specificity but shares a large set of user needs. These require a common set of design solutions which have been "captured" by the MUSE application framework. To refine and empower the framework, we extended it with a number of features which we abstracted from the analysis of a large number of existing Web applications for cultural tourism.

The MUSE design framework is described using the concepts and notations of the W2000 conceptual design model. W2000, developed within the EU funded project IST - 2000 – 25131 "UWA" (Ubiquitous Web Applications), provides the modeling primitives to specify, in an implementation independent way, the various aspects of a multichannel Web application. W2000 distinguishes among the following design tasks (Garzotto et al, 1993, Garzotto et al. 1995, Baresi et al., 2000, Baresi et al. 2001):

- *Information design*: it describes the content delivered by the application. The result of information design is a set of *information design schemas*.
- *Navigation design*: it defines how contents are structured for the purpose of navigation. The result of navigation design is a set *navigation design schemas*.
- *Interaction design*: it defines the user interaction options and the dynamic behavior of the application at run time, in response to the user actions
- *Presentation design*: it defines how navigation structures and interaction elements are presented from the delivery device. The result of presentation design is a set of *presentation design schemas* and *screen templates*.

The notation adopted in W2000 to specify the various schemas is an extension to the Unified Modeling Language UML (Booch 1998, Conallen 1999). UML is a software engineering de-facto standard that provides an extendible kit of graphical elements to model, in an object oriented fashion, data structures and behavioral properties of any software application.

A design activity which is orthogonal to the above-mentioned tasks is *customization design*. During customization design, the designer defines which application features – content, navigation, interaction, presentation – need to be specialized to the context. The term "context", in W2000, comprises all aspects concerning the situation of use: device characteristics, user profile, physical and geographical position, etc. W2000 also distinguishes between *customization at design time* and *customization at run time*. Design time customization means that the designer specifies different design schemas to address the needs of different contexts. Run time customization means that a set of adaptation rules define how some detailed design solutions are dynamically adapted to changes of context.

In MC applications for cultural tourism, both design-time customization and run-time customization are needed. Design time customization is needed to specify the content, navigation, and interaction features of the application on each different channel. Context-aware behaviour can be better expressed by using adaptation rules that define which elements must be presented on the current channel in response to user position or orientation changes.

Although MUSE framework provides general patterns for run-time adaptation rules, these are not discussed in this paper for lack of space. In the rest of this section, we will discuss issues related to design-time customization.

The problem of customizing the design to the requirements of a specific channel can be addressed in several ways:

Approach 1: Customization starts during content design.

Different information design schemas are produced, each one defining the content structures that are specific for each channel. A channel specific navigation schema and presentation schema are defined for each information schema.

Approach 2: Customization is postponed to navigation design.

The designer specifies a single information schema which defines all possible content structures for all channels. During navigation design, these are filtered and restructured for each specific channel, resulting into a set of channel-specific navigation schemas coupled with the corresponding presentation schema. In data base language, we can say that these navigation schemas define *channel specific views* on the common pool of information defined by the information schema.

Approach 3: Customization is postponed to presentation design.

In this case, the designers produces a single information schema, a single navigation schema, and multiple presentation schemas. The specification of channel specific contents and links is described in the presentation structures.

Each approach has different advantages and disadvantages. Approaches 2 and 3 are more effective for authors and implementers. Since the information schema is the "summa" of all information needed for all channels, it can be regarded as an authoring check list, mentioning all content that authors must produce. For implementers, designing the "pool of contents" (i.e., the shared data base upon which the architecture of a multichannel application is built) from a single information schema is easier and less error-prone than designing it from multiple information schemas; for example, the probability of introducing information redundancies and inconsistencies is minimized. The disadvantage of approach 2 is that implementers must look at the navigation schemas in order to identify which content will be delivered to which channels. With approach 3, the implementer must inspect the page structures of each channel-specific presentation schema in order to define the content and link views that are available in each channel.

MUSE has adopted approach 2 (multiple navigation schemas); this seems to be the most effective compromise. In the following sections, we will discuss the general information schema of the MUSE framework, and we will sketch some portions of some channel specific navigation schema. Presentation schemas are omitted for lack of space, and illustrated via some screen examples. Finally, we will discuss some aspects of the interaction design.

MUSE Information Design Schema

The MUSE Information Schema includes the main classes of multimedia information elements that satisfy the information requirements of an MC cultural tourism application, abstracting from the requirements of each specific channel. In principle, "any" MC application for cultural tourism can reuse this schema as a design skeleton which can then be filtered and integrated according to the needs of a specific domain.

According to W2000, two main categories of content structures are described in the information schema: *hyperbase layer structures* and *access layer structures*.

Hyperbase layer structures define the core content sources of an application. These are described

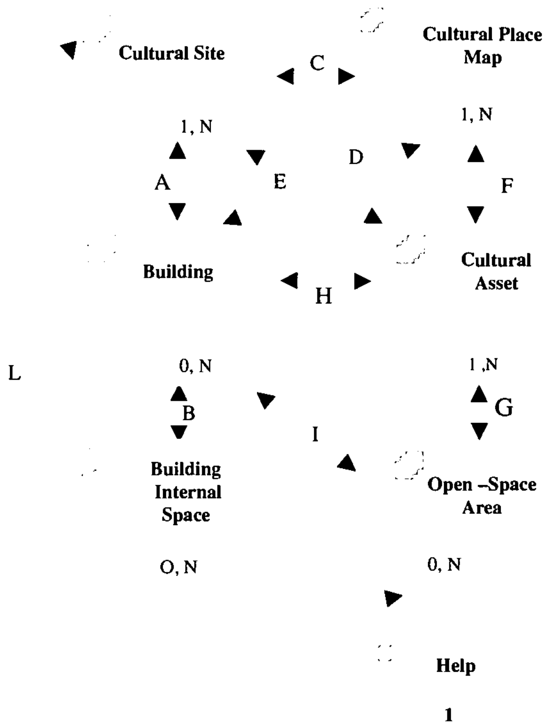

Figure 3: The MUSE Hyperbase Information schema - in-the-large view

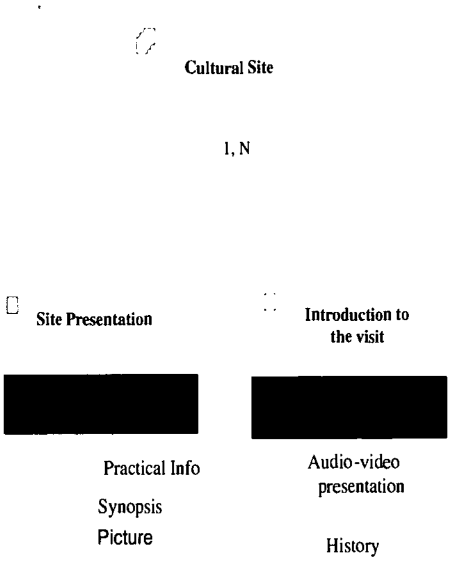

Figure 4a: The content of Entity Type "Cultural Site"

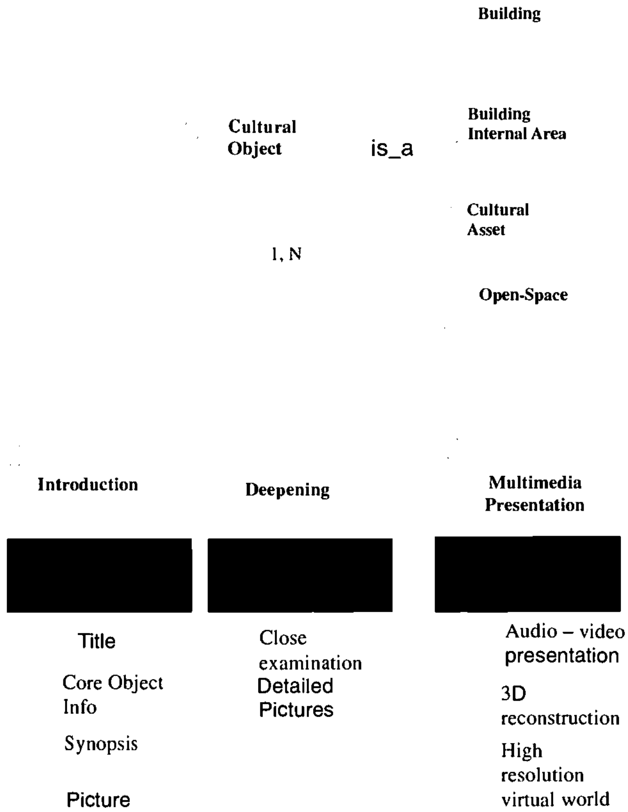

Figure 4b: The content of Entity Type "Cultural Object"

in terms of *entity types* and *semantic association types*. Entity types model the main "objects" of the application domain as classes of information structures called "entities". Semantic association types model the relevant relationships among these entities.

Figure 3 shows that the relevant classes of information objects deal with Cultural Sites (the tourism places), their Cultural Assets, Buildings, Internal Building Spaces and Open-space areas. In addition, the application may include Maps (more or less sophisticated and rich of multimedia content) and some Help Information.

Figures 4a and 4b describe the content elements associated to the various entity types (Map Specification is omitted). Content elements within entity types are grouped in sub-structures called Component Types, characterized by a black banner. Except Cultural Place and Map, all entity types share the same content structure. In the object-oriented fashion of UML, this can be specified by defining an *Abstract Entity Type* - called *Object of Interest* in our case – and by saying that entity types Building, Cultural Asset, Building Internal Space, Open-Space Area are all sub-classes of Object of Interest (see figure 4b).

Hyperbase structures define the content "building blocks" of an MC application and classify them according to their intrinsic meaning and nature. Access Layer structure defines groups of hyperbase elements, called *Collections*, which provide complementary ways of organizing the contents. For example, a collection called "Highlights" is defined to group the most appealing Cultural Objects for a short visit of a cultural site. The goal of access layer structures is to drive the user towards the core application content that is represented by the hyperbase structures: During navigation, access layer structures are traversed before and in order to access the hyperbase.

According to the W2000 design model, a collection is described by its *members* and by a *center*. The center is an information element which describes the collection itself, e.g., by listing its members and by providing a short introduction to the collection subject. Collections can group hyperbase elements or other collections. In the latter case they are called *nested collections*.

In the MUSE application framework, the access layer structure includes the following collections:

- A collection for each entity type, which groups all entities of a given type ("Buildings", "Cultural Assets", Open Spaces", "Internal Spaces") . The center of these "By-type" collections includes a list of member descriptors (e.g. titles and miniaturized picture).

- A collection for each relevant theme, which groups the main entities relevant for a given topic. The center of a thematic collection includes a list of member descriptors (e.g. titles and miniaturized picture), and a short introduction to the collection theme.

- The collection "Highlights", which provides a selection of the most important elements for a short visit. Its center includes a short introduction to the visit, and the list of member descriptors which can be eventually visualized on a map showing where the various elements can be found on the cultural site.

- The collections "Geographic Tour", which groups "Cultural Objects" according to spatial criteria (e.g., in a large archeological site like Pompei, all monuments and assets "around the forum")

- The Collection "Album", which is created by the user while he or she is using the application as discussed in the memories channel.

- The nested collection "Index of Themes", which groups all the thematic collections

- The nested collection "Index of Geographic Tours", which groups all the collections' "Geographic tours"

- The nested collection "Index of Topics", which groups all the "By-Type" collections

MUSE Navigation Design

The MUSE Information Schema only describe *which contents and groupings* are relevant for all possible channels of an MC cultural application. It does not provide any specification of which contents are navigated during the use of the application in a specific channel, nor how they are navigated. These aspects are described by the MUSE Navigation Schemas.

Each channel-specific navigation schema specifies a set of *node types* and *link types*. A node is an atomic unit of navigation and presentation in a specific channel. This means that navigation links connect nodes (and not node portions), and the effect of navigation is to activate an entire node: all content items in this node are loaded and made available on the page together (possibly together with other nodes)

During navigation design, the designer should

1. select the information elements in the information schema that are made available in each specific channel;

2. group them into nodes;

3. define the links among nodes that allow the user to exploit the semantic associations and the collections defined in the information schema.

The following diagrams show some portions of the MUSE framework navigation schemas for the different MUSE channels. Fat black arrows identify the entry to some sets of nodes and links; checkmarks identify "landmark" nodes, which can be reached from any node in the hyperspace.

These diagrams shows that the two on-site channels focus on multimedia content which has a more direct, emotional impact on the user and is a more appealing complement to what the user can experience on site. When the user employs these devices, multimedia contents are the first ones to be presented - audio-visual content for the mobile channel, and high resolution, interactive 3D content for the on-site stationary channel. Further (mainly based on static media) deepening can be retrieved on demand. In contrast, the navigation flow on the conventional Web channel starts from text and pictures, and proceeds to multimedia if the user wants further deepening.

Muse Interaction Design

Interaction design defines how the users interact with the MC application in each different channel, and how the application behaves in response to user actions.

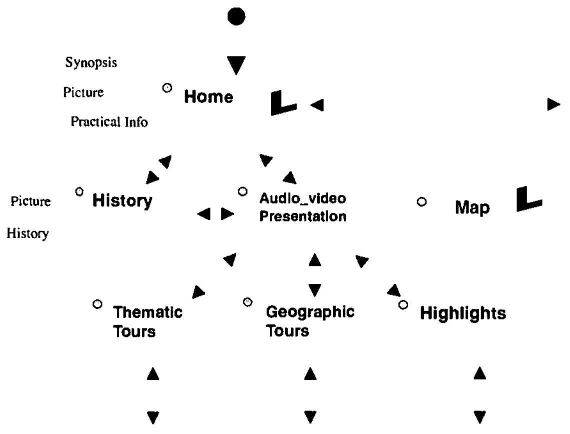

Figure 5a: Conventional Web Navigation - High-level Access Structures Navigation Schema

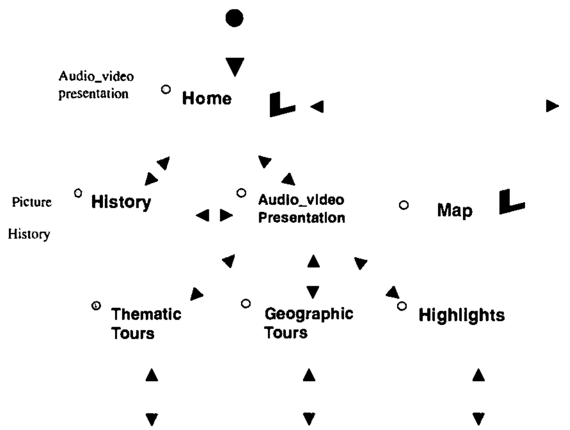

Figure 5b: On-site Mobile Channel - High-level Access Structures Navigation Schema

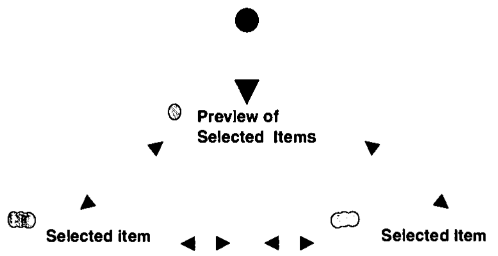

Figure 5c: Memories Channel - Navigation Schema

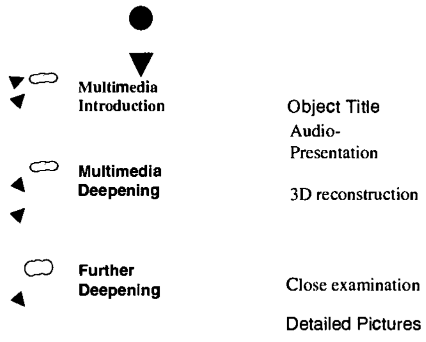

Figure 6a: On-site Mobile Channel – Cultural Object Navigation Schema

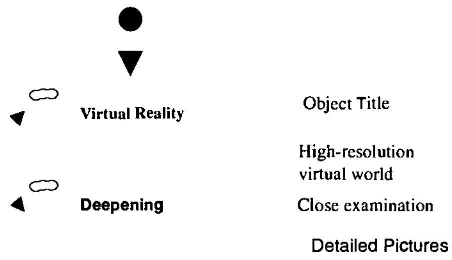

Figure 6b: On-site Stationary Channel – Cultural Object Navigation Schema

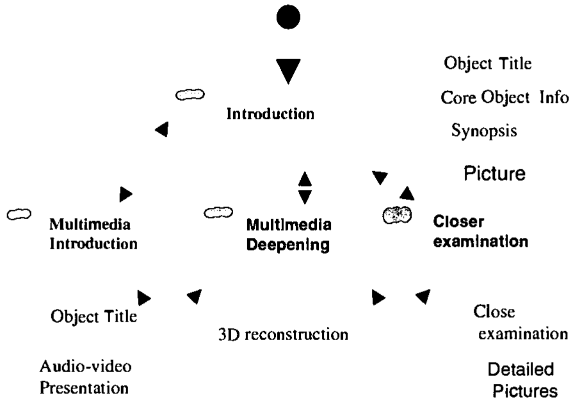

Figure 6c: Conventional Web Channel – Cultural Object Navigation Schema

Interaction in the conventional Web channel

The interaction mode is *free browsing*, i.e., standard hypertextual navigation: user can activate any link defined in the navigation model, and the effect of link execution is to load and display the link destination node.

Interaction on the on-site mobile channel

MUSE supports three interaction modes: *free browsing, context-aware navigation,* and *memories-building*. Context aware navigation is the default, but the user can change interaction mode at any point during a session of use.

Free browsing

This is standard hypertextual navigation as defined for the conventional Web channel. The only difference is that the on-site mobile channel provides different content and navigation structures, as defined by the corresponding navigation schema. In free browsing, context-awareness capabilities of the device are off. What is presented on the device is totally under the user control and is not affected by his/her current position in the physical space.

Context-aware navigation

This interaction mode exploits the device capability of detecting the current geographical position of the user and of delivering context-dependent content. MUSE provides different forms of context-aware navigation which offer different levels of visibility of the content and different levels of control by the user:

Pure context-aware navigation

In this interaction mode, the application has the complete control. When the user reaches a place or an object in the tourist site associated to some audio-visual content, the device automatically delivers this information on the mobile device. It plays until its end, or until the user reaches a different position which is associated to a different content. The users are largely passive. They can only control the state of multimedia elements (e.g., suspending or resuming or stopping the audio-visual presentation) or switch to a different interaction mode.

Guided context-aware navigation

This interaction mode enhances pure context aware navigation with the notion of guided navigation. The application guides the users along a path on the site they are visiting, pinpointing the relevant objects there. These objects correspond to those grouped by the collections "Highlights", "Thematic Tours", "Geographical Tours". At any point of time, what is presented to the user is calculated by the application on the basis of both the user current position in the physical space and the currently active guided tour. The interaction metaphor is the vehicle satellite navigator: after the users have selected a tour of interest, the application highlights the path on an active map and shows the users' current positions. If the users reach a place or an object, either inside or outside the current path, which is associated to some audio-visual content, the device automatically loads and (dis)plays it. The application can detect any user deviation from the current tour path; wandering causes a sound warning and the display on the map of the shortest path to reach the currently active tour. In guided context-aware mode, the users are slightly less passive than in pure context-aware navigation. Beside controlling the state of multimedia elements (e.g., suspending or resuming or stopping the audio-visual presentation) and changing interaction mode, the users can ask for the map display, can change tour by returning to the tour selection menu, or can continue the exploration on a different tour. The latter possibility is available only when the users reach a place at the intersection of different tours. These tours are shown on the map and the users can choose any of them (the selection becomes the currently active tour).

Integrated context-aware navigation.

It merges the capabilities of free navigation and context aware navigation. After the application displays some content in response to a change of user position or orientation, the users are allowed to explore related content and to browse as they can do in free navigation mode. Since the context-aware capability remains on, any physical movement of the users may cause a position-dependent replacement of the content which is currently active on the device.

Memories-building mode

The memories-building mode provides the users with the operations needed to manipulate the content for the CD memories channel. At any point during the use of the mobile device, the users can execute any of the following operations:

• *mark* the currently displayed screen or *take a picture* of what they are looking at (using a digital camera incorporated in Whyere). The effect of these operation is to update the memory album

• *view* and *navigate* the memories album

• *modify* the memory album by deleting some selected items, re-ordering them, or including comments and annotations

Any of the above operations activates the memories-building mode and suspends the current interaction mode, which can be resumed after the completion of the operations.

On-site stationary channel

In any place where the on-site stationary channel (a large screen controlled by a powerful graphic station) is available, users can switch to this channel and interact with it. What appears on the large screen is not a zoomed display of what is presented on the mobile device, but complementary high-resolution, interactive multimedia content which is related to the place where the stationary channel is located. Using the on-site mobile device as a control and pointing device, users can operate on the state of multimedia elements, explore the 3D space, and traverse navigational links. The content space available for navigation is a view of the whole content designed for the on-line stationary channel, filtered according to the position of the stationary device.

Muse Presentation Design

The goal of presentation design is to define how contents and interaction elements are displayed on the different delivery channels. This activity comprises two sub-tasks: defining, for each channel, the *conceptual presentation* and the *concrete presentation*. Conceptual presentation defines *which* content and interaction elements are displayed simultaneously on the screen. Concrete presentation defines their lay-out graphical properties - space allocation, colors, typographical properties, images frames, etc..

The MUSE presentation framework mainly addresses conceptual presentation which can be reasonably shared and reused by a large number of applications in the tourism domain. Concrete presentation is largely "application specific" and may take into account the general visual communication strategy of the client institution.

The MUSE conceptual presentation framework comprises a presentation schema for each different channel. Each presentation schema defines a set of *conceptual page templates* which specify the types of nodes and links that must appear on the screen of a specific channel (Strothotte 1998). For lack of space, in this section we will focus on the presentation design for the on-site mobile channel only.

Given the limited size of the mobile device screen Whyre (6.4 inches), each page template includes a single node type and a placeholder for all link types outgoing from that node type in the navigation schema of the on-site mobile channel.Another constraining aspects of the on-site mobile channel which has an impact on presentation design is the lack of any pointing device, like a mouse, a trackball or a hand touch, to allow users to directly select any point on the screen. To interact with the system (e.g., to activate a link), the users must rely upon *physical buttons* available on the screen plastic frame. This factor introduces some intriguing requirements for concrete presentation: it forces the presentation designer to have a limited number of interaction elements on the screen and to place them in a position which allows users to "select" them using physical buttons only.

The Whyre "keyboard" includes four hardwired buttons and six "multimodal" buttons.The hardwired keys support the following functions: *undo, SOS (help), orientation, take a picture*.

The multimodal buttons are symmetrically located on both sides of the device. In this way, they are easily reachable by the users. As shown in the following pictures, each interaction element is placed on the screen close to a multimodal button which

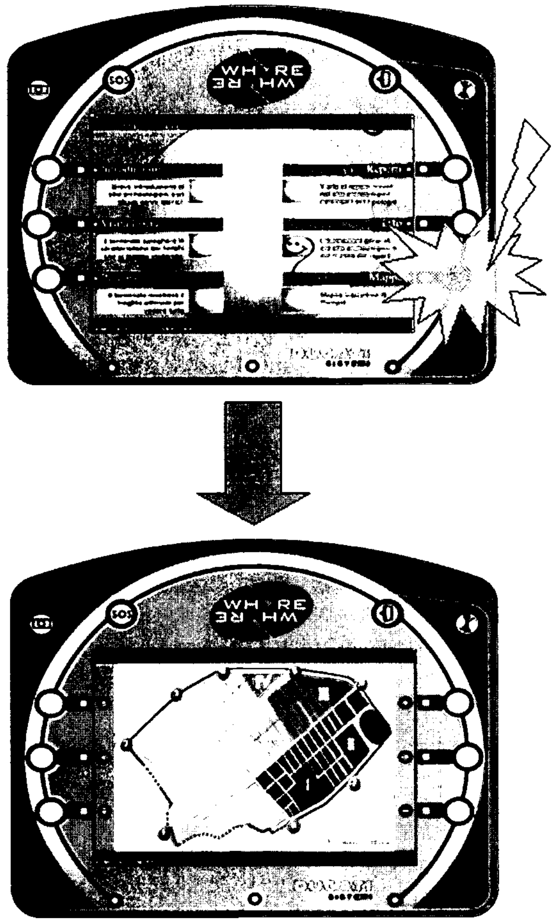

Figure 7: From the virtual entrance of Pompei to the site map

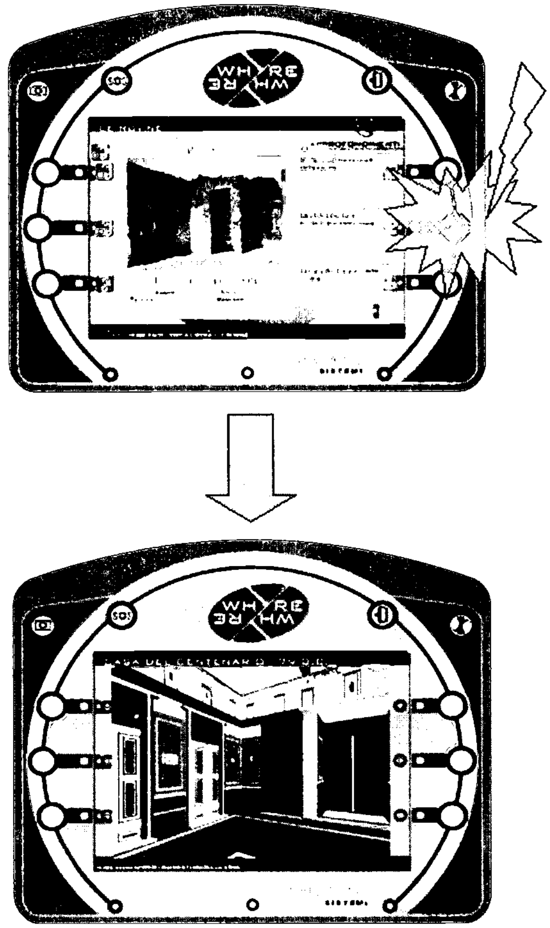

Fig.8: Entering a virtual reconstruction of La Casa del Centenario while navigating in integrated context-aware mode

can be used to select that interaction element. Pressing the button executes the closest interaction element. The semantics of a multimodal button (which is the reason for this name) depends on the semantics of the closest interaction element.

The first screen displayed in figure 7 is the entry point ("home") of the Pompei application on the on-site mobile channel. The following interaction choices are available:

- Get an introductory presentation on the site
- Show the site map
- Select a navigation mode (one key for each mode)
- Get help on the current screen

By pressing the closest button to the "Site map label" (see left image – last button on the right side), the map of the archeological site, organized in different "insulas", is displayed (see right image).

In Figure 8 we see the effects of pressing a different button (the mid right button in the left image). We suppose that the users are in the integrated context aware mode. When they enter *La Casa del Centenario* in Pompei, their position is detected by the application: an image showing the house atrium automatically appears (see left side image), and a sound comment starts. The users can choose a deeper description by pressing the mid left button, and can virtually enter the *atrium* as it was before the 79 A.C. Vesuvius eruption, according to the re-

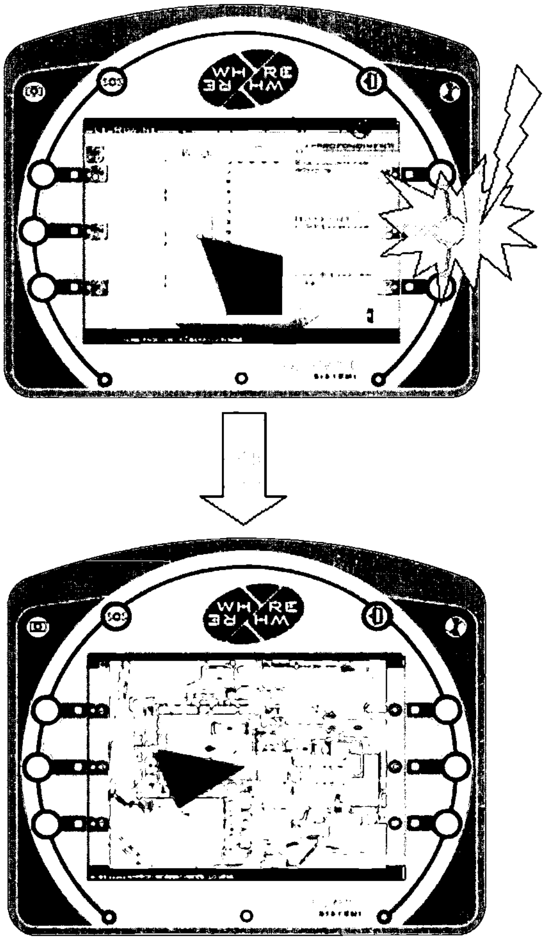

Fig 9: Displaying a guided tour path in guided context aware mode

construction validated by the archaeologists (see the right side image) (Scagliarini, 2001).

Figure 9 provides an example of the multimodal behaviour of buttons and of guided context aware navigation. After the orientation hardwired button is pressed, the application displays the site map showing the visitors' position and orientation, as well as the list of the reachable places belonging to the current guided tour (left side image). If a desired place is selected by pressing the mid right button (see left image) the application displays a zoomed map (see right image), showing the area included between the visitors' current location and the selected section, and the recommended connecting path.

The Muse Interface Implementation Framework

The Muse interface implementation framework is called *MuseXP*. It consists of a pair of interface definition tools (an editor and a compiler) coupled to an interface-independent execution engine (called *MuseXP Runner*). With *MuseXP*, MC systems interface design is approached at a higher level of abstraction with respect to standard Web applications. *MuseXP* targets are channel-specific interfaces in the cultural heritage domain. The framework is designed to be reconfigurable in order to be reused in other domains, should the opportunity arise.

Within *MuseXP* each channel-specific navigation and interaction design is mapped on to an XML file called *ch_interface.xml*. This file is a centralized interface control point easily handled by its creator and easily interpreted by the interface independent execution engine. XML (standing for *eXtensible Markup Language*) was chosen for its extensibility: XML syntax was programmed with a number of tags finalized to the MuseMC system semantics. Interface presentation design is implemented filling in device-dependent HTML pages with the screen layouts and the graphic objects.

MuseXP Interface design workflow

Given the conceptual design of a specific interface within a *MCMuse* system, how is the operational design carried out? How is the interface implemented on the target device?

In terms of the *MuseXP* architecture, these questions need to be read as follows: how are the *ch_interface.xml* file and the screen graphic layouts created? How are they processed on the target device by the *MuseXPRunner*?

Fig. 10 is a block diagram of the *MuseXP* system currently under development within the MUSE project.

Within the envisaged framework, the interface operational design is split into three steps:

1. Navigation and interaction specification

2. HTML containers generation

3. Interface graphic design

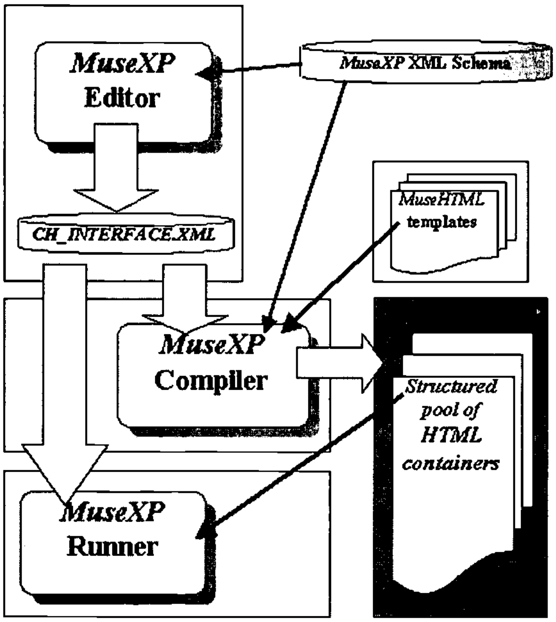

Figure 10: Overview of MuseXP workflow

The first activity is expected to be carried out using a MUSE specific visual tool called *MuseXP editor* currently under development. *Muse XP editor* maps the interface navigation and interaction designs into the text file *ch_interface.xml*. Since the editor is not yet available, *ch_interface.xml* has to be generated using an uncommitted text editor.

The second step is carried out automatically by an off-line tool called *MuseXP Compiler.*

The *compiler* generates a tree of all HTML pages required to implement the interface. These pages are called "*HTML containers*" because at this stage they only include the qualified structure-definition objects (i.e. the references to the required viewers and their attributes), but they do not include any graphic and metric layout information for the page presentation. The "HTML containers" pool is the set of all presentation instances, hierarchically arranged as declared in *ch_interface.xml*.There are no hypertextual links among these "Web pages". Navigation within the pool will be performed by the execution engine according to *ch_interface.xml* specifications. Each page is externally controlled according to the DHTML extension *DOM (Document Object Model)*. In this way the contents to be delivered may be specified at run time.

The third step consists of the actual presentation implementation. During this step the graphic objects used throughout the interface are created and the *HTML containers* are filled up with the missing presentation items.

This activity can be carried out by Web interface designers or persons with similar skills, using standard graphics tools as well as HTML visual editors.

Once the HTML containers are graphically filled, they are saved on the target device with the *ch_interface.xml:* together, they specify the run-time rules-of-behaviour to *MuseXPRunner,* the execution engine that implements the designed interface.

The proposed framework has many advantages with respect to interface specific implementations:

- run time specification of both contents and page layout make a channel-specific interface reusable in many applications of the same domain

- an interface independent execution engine cancels programming from the list of design and implementation activities

- Separating structure from graphics is an additional answer to the complexity handling problem: it reduces the skills required for the most time-consuming tasks and it makes for easier design management (graphic layout, for example, can be modified without having to recompile the interface structure)

Conclusions

State-of-the-art Web technologies offer the opportunity to deliver interactive multimedia content on-demand. Similarly, hardware technologies offer a wider and wider variety of devices, providing ergonomy and interaction modes much better customized to the users' context and location than a standard PC. These devices range from wireless connected mobile terminals exhibiting small and bright screens to large plasma display providing nearly immersive fruition in dedicated areas. The combination of these technologies within the same system offers the potential to provide ubiquitous services originating from the same information base and properly optimized to the fruition conditions. The user interfaces in such multichannel systems are different from standard Web clients. Web technology may be used, but the application design should be approached at a higher level of abstraction. Within the project MUSE this multichannel paradigm is currently under investigation. A framework supporting the design and implementation of heterogeneous interfaces in multichannel environments was developed, and its use in cultural tourism applications has been described in this paper. The framework was used to implement mobile and stationary on-site and off-site channels for three MC systems addressing three cultural tourism "case studies": the *insula del Centenario* in the archaeological area of Pompei, the *History of Science Museum* in Florence, and a section of la *Certosa di San Martin,o* the outstanding historical Museum of the city of Napoli. The MUSE system will be demonstrated at these locations later this year, as soon as Whyre, the interactive multimedia terminal specially designed for on-site fruition in cultural heritage establishments, is ready.

Acknowledgements

MUSE is funded by the Italian government (MIUR, the Italian Ministry for University and Research), within the framework of the National Research Program on Cultural Heritage Parnaso. The project is conducted by Ducati Sistemi S.p.A. in cooperation with Cineca, Politecnico di Milano, Sinet and the University of Bologna. Content providers as well as hosts for extended system concept verification are *Istituto e Museo di Storia della Scienza* (Florence), *Soprintendenza Speciale del Polo Museale di Napoli* and *Soprintendenza Archeologica di Pompei*. The authors wish to thank Prof. Pietro Giovanni Guzzo for establishing the connection between MUSE and the project "Pompei - Insula del Centenario"; Prof. Nicola Spinosa and his wonderful team for their continuous and irreplaceable support; Prof. Paolo Galluzzi for his conceptual contribution to the fruition model; Prof Francesca Bocchi for managing the training program associated to the research project; Ravi Nagaraj and Wen-Hann Wang with Intel Corporation for fruitful support in the design of Whyre hardware; Prof. Daniela Scagliarini and Antonella Guidazzoli with their teams for the development of the 3-D models of the *Casa del Centenario*.

References

Abowd, G. D. (1999). Software Engineering Issues for Ubiquitous Computing. *International Conference on Software Engineering (ICSE)*, Los Angeles.

Baresi L., F. Garzotto, P. Paolini, (October, 2000). From Web Sites to Web Applications: New Issues for Conceptual Modeling. *WWW Conceptual Modeling Conference*, Proceedings, Salt Lake City.

Baresi L., F. Garzotto, P. Paolini, (January, 2001). Extending UML for Modeling Web Applications. In the *34th Hawaii International Conference on System Sciences*, P,roceedings Maui (USA).

Bocchi F. et al. (August, 1999). The 4D Virtual Museum of the City of Bologna, Italy, *Siggraph 1999*, Los Angeles.

Booch G., I. Jacobson and J. Rumbaugh (1998). *The Unified Modeling Language User Guide*. The Addison-Wesley Object Technology Series.

Conallen J. (Oct. 1999). Modeling Web Application Architectures with UML. *Communications of the ACM*, 42:10, 63-70.

Fayad M.E. and D.C. Schmidt, (Oct. 1997). Object Oriented Application Frameworks in *Communications of the ACM*, vol. 40, n. 10.

Garzotto F., P. Paolini, and D. Schwabe, (January, 1993). HDM - A Model-Based Approach to Hypertext Application Design. *ACM Transactions on Information Systems*, Vol. 11, No. 1.

Garzotto F., P. Paolini, L. Mainetti, (Aug. 1995). Hypermedia Design, Analysis, and Evaluation Issues. *In Communications of the ACM*, Vol. 38, N. 8.

Garzotto F., P. Paolini, D. Bolchini, and S.Valenti, (15 Nov. 1999). "Modeling by patterns" of Web Applications, in *Proc. Of the International workshop on the World-Wide Web and Conceptual Modelling*, WWWCM'99, Paris, 293-306.

Guidazzoli A. et al., (July, 2000). Tailored Virtual Tours in Cultural Heritage Worlds *Proceedings Siggraph 2000*, New Orleans.

Johnson R.E., (Oct. 1997). "Frameworks= Components + Patterns", in *Communications of the ACM*, vol. 40, n. 10.

Oppermann R., and M. Specht, (September, 1999). A Nomadic Information System for Adaptive Exhibition Guidance, *Proc. of the International Conference on Hypermedia and Interactivity in Museums (ICHIM)*, D. Bearman and J. Trant (eds.), Washington.

Ozen, B.T., Kilic, O., Altinel, M., Dogac, A., (May, 2001). Highly Personalized Information Delivery to Mobile Clients, *Proc. of 2nd ACM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDE '01)*, Santa Barbara, USA.

Perkowitz, M. and O. Etzioni, (May, 1999). Towards Adaptive Web Sites: Conceptual Framework and Case Study, *Proc. of the The Eighth International World Wide Web Conference, (WWW8)*, Toronto, Canada.

Salmon Cinotti, T. et al., (Milano, 2001). MUSE an Integrated System for Mobile Fruition and Site Management. In *ICHIM01 (International Cultural Heritage Informatics Meetings*; Milano, Politecnico, Sep 3-7 2001), Proceedings, Vol. I, 609-621.

Schmid, H.A., (Oct. 1997). "Systematic Framework Design", in *Communications of the ACM*, vol. 40, n. 10.

Scagliarini, D. et al., (November, 2001). Exciting understanding in Pompeii through on-site parallel interaction with dual time virtual models, *VAST01*, Glyfada (Greece).

Schwabe, D., G. Rossi, L. Emeraldo, F. Lyardet, (1999). Web Design Frameworks: An approach to inprove reuse in Web Applications. *Proceedings WWW99 Web Engineering Workshop*, Springer Verlag.

Strothotte, T., (1998). Computational Visualization - Graphics, abstraction and interactivity, Springer-Verlag.

Weiser, M., (July, 1993). "Some computer science issues in ubiquitous computing", *CACM*, Vol. 36, No. 7.

About the Authors

Authors' Biographies

Ivana Alfaro holds a Master's Degree in International Economic Policy and Market Development from the American University in Washington D.C. She joined the Cognitive and Communications Technology Division at ITC-irst in November 2000, where she conducts research in mobile multimedia guides for museums. Her interests include the study of interface usability analysis and evaluation as well as market research for future technology applications for multimedia guides.

Pilar Almeida is a Brazilian national currently completing a PhD at the Graduate School of Human Informatics, University of Nagoya, Japan. He acquired his Masters in the same department from 1998 to 2000. The subject of his research was the possible use of interactive narrative for online museums.

Steve Allison-Bunnell has worked as a writer and producer in New Media since 1995, when he was the founding Nature and Science Editor of the Discovery Channel Online. He has written and developed non-fiction educational materials for Discovery.com, Britannica.com, the USDA Forest Service, the JASON Project, and the Shedd Aquarium. Before turning to online multimedia, he wrote and produced science and environmental reporting for television, public radio, and print. He holds a PhD in Science and Technology Studies from Cornell University, and was a Smithsonian Institution Predoctoral Fellow.

Patricia Barbanell holds a Doctorate from Columbia University and has extensive experience in both K-12 education and Museum education. She has worked for over 20 years developing integrated programs that serve both museums and schools. Her specialties are integrated arts, multicultural programming and technology integration. She has presented dozens of presentations at professional conferences and has published several papers in professional journals. She is past president of NYS Art Teachers and NY Council of Educational Associations and helped to write the NYS Learning Standards for the Arts.

Lorrie Beaumont specializes in the evaluation of informal and formal education programs. Lorrie spent four years as an educator and evaluator at the DuPage Children's Museum prior to becoming an associate with Selinda Research Associates as well as an independent evaluation consultant. She has advised several museums on teaching techniques for educators. Her research interests include working with young children and families in informal settings, training of educators in museum settings, and professional development for teachers. Lorrie is currently a doctoral student in Educational Psychology with an emphasis in Program Evaluation and Child Development.

Laura Matthias Bendoly has served as Curriculum and Distance Learning Manager in the education department of the Atlanta History Center since March, 2002. Her responsibilities there include the research and writing of education materials for the K-12 audience, the development and delivery of online resources (social studies lesson plans, interactive games and classroom activities), and the creation of distance learning programs (live, teleconference performances for the K-12 grades). Previous education positions include Coordinator of Academic Services at the High Museum of Art (Atlanta), and Curriculum-Structured Tour Manager, at the Indiana University Art Museum. Her degrees include an MA in art history from Indiana University, an MFA in creative writing from Western Michigan University, and a BA in English and French from the University of Notre Dame.

Stefania Boiano has been working since 1995 for Web agencies and major Italian Science Museums in Naples and Milan as Web project manager and art director. Now she is a freelance consultant.

Lee Anne Burrough is responsible for coordinating development and implementation of the award winning Kids Design Network internet-based program at the DuPage Children's Museum. She has a B.S. and M.S. in Geology and teaches Earth science and geology at the college level. Independently, she combines her interests in science, children and education to develop and present Girl Scout, elementary and pre-school hands-on science programs including paleontology, geologic history, energy and hazardous waste. Prior to becoming involved in education, she investigated and remediated hazardous waste sites.

Ethy Cannon has twenty years of experience programming educational software and Web sites. As a programmer for Eduweb, she has created server-side and client side programs for many projects, including several award winning projects: Lights, Puppets, Action for The Children's Museum of Indianapolis and Art Tales for the National Museum of Wildlife Art. Previously, she was a senior software engineer at MECC and The Learning Company, where she worked on such educational projects as Oregon Trail, Yukon Trail, Big Science Ideas, Compton's Encyclopedia On-Line and Oregon Trail On-Line. She has an M.A. and a B.S. in computer science and an M.M. in cello performance and music literature.

Morgana Caldarini is the president of Jargon, a Milanese Web agency specialized in advanced-technology projects. www.jargon.it

Nicoletta Di Blas graduated in Classics and obtained a PhD in Linguistic Sciences from the Catholic University of Milan. She currently teaches Sociology of Communication at the Politecnico di Milano. Her research interests focus on linguistic themes, on the relation between rhetoric and new technologies, on usability and educational applications, in particular as regards cultural heritage. She played a major role in the organization of ICHIM2001. She coordinates the cultural content aspects of the SEE project for Politecnico di Milano, in cooperation with the curator of the Shrine of the Book.

Steve Dietz is Curator of New Media at the Walker Art Center in Minneapolis, Minnesota, USA, where he founded the New Media Initiatives department in 1996. He is responsible for museum informatics and programming the online Gallery 9, including more than 20 net art commissions and one of the earliest archive-collections of net art, the Walker's Digital Arts Study Collection.

John Falco has served as the Superintendent of the 9,000 student urban school system of Schenectady NY for 3 years and was Deputy Superintendent for 6 years prior to that. In addition to Project VIEW, Dr. Falco has led the development of a number of transformational projects in the Schenectady City Schools District, like the Capital Region Science Education Partnership (CRSEP), a multi-district Local Systemic Change Initiative Project funded by the National Science Foundation. Dr. Falco was recently named the NYSCATE Superintendent of the Year for his outstanding achievements in integrating technology across the Schenectady City School District. Dr. Falco holds a Doctorate in Educational Administration from Seton Hall University, where his research centered on improving reading skills for struggling emergent readers.

Giuliano Gaia worked for five years as Web project manager for the Science Museum of Milan. Now he is a freelance consultant and runs the discussion mailing list Musei-it www.musei-it.net.

Franca Garzotto is Associate Professor of Computer Engineering at the Department of Electronics and Information, Politecnico di Milano, where she teaches classes of Fundamentals of Computing, Hypermedia Design, and Multimedia Systems. She has a Degree in Mathematics from the University of Padova (Italy) and a Ph.D. in Computer Engineering from Politecnico di Milano. Her research has focused on document modeling, hypermedia design, hypermedia usability, mainly applied to e-learning, cultural heritage, and e-commerce. She has been involved in various European Commission R&D projects in the above fields. She was tutorial chair and/or member of the technical program committee of several editions of many international conferences:

Stefan Göbel heads the department Digital Storytelling at the Computer Graphics Center in Darmstadt, Germany. He studied computer sciences at the technical university of Darmstadt and finished his thesis at 1997 in the field of geographic information systems. After his thesis, he has been working at the Fraunhofer Institute for Computer Graphics and is the author of numerous publications within the research field of GIS, user interaction and information visualization.

Steve Guynup has been a leader in web3D development for the past eight years. He has presented virtual works at SIGGRAPH (Web3D Round-up and the Art Gallery in 1998, 1999 and 2000), VRML / Web3D Conferences (1997, 1998, 1999) won design awards from The Contact Consortium and Blaxxun Interactive and presented projects at the Atlanta Arts Festival (1997) and Digital Americana (1998, Orlando Museum of Art) He was also a senior artist with GRA Interactive and as a web3D interface designer for 1996 Ars Electronica winner, Andy

Best's MeetFactory. He currently a graduate student at the Georgia Institute of Technology and is looking to continue his exploration of the virtual as professor of digital media

Kate Haley Goldman has worked on evaluation projects with the National Aquarium in Baltimore, Disney's Animal Kingdom, Atlanta History Center, Astronomy Society of the Pacific and the Cleveland Museum of Art. She has a bachelor's degree in Anthropology from Bryn Mawr College and has extensive training in educational measurement, statistics and evaluation in learning and technology. Previously Kate worked in several departments of the U.S. Holocaust Memorial Museum, primarily concentrating on audience research. Her research priorities include the long term impact of museum visits and investigation of free-choice learning in new media environments.

Susan Hazan is the Curator of New Media Education Unit, Head of the Internet Office at The Israel Museum. Her Master is in Media and Communications, Goldsmiths, University London, UK. Her doctoral research focuses on an exploration of electronic architectures in the contemporary museum and New Media in education, a field in which she has published many papers. She was keynote speaker at the Museum and the Web conference, 2001 and has served on the program committee for Museums and the Web and ICHIM for several years. She coordinates the major communication-interaction aspects of the SEE project for the Israel Museum.

Caro Howell established Tate Modern's outreach education programme in the three years prior to the gallery opening and originated the first peer-led youth programme at a London museum. She is currently the Curator for Special Projects at Tate Modern developing programmes, facilities and training for people with disabilities, particularly sensory impairments. She is also a freelance access consultant and gallery educator.

Jane Hunter is a Distinguished Research Fellow at the Distributed Systems Technology Centre, at the University of Queensland. Her research interests are multimedia metadata modelling and interoperability between metadata standards across domains and media types. In 2002 she undertook a Queensland-Smithsonian fellowship at the National Museum of the American Indian CRC investigating the application of multimedia metadata and rights management standards to the development of online collections of indigenous resources. This work is continuing through a number of collaborative projects with indigenous communities in Australia and the US.

Ido Iurgel holds a masters degree in philosophy, social psychology and general linguistics, and a German diploma (the equivalent of an M. Sc.) in computer science. Formerly he was a scientist at the Ruhr-University of Bochum in Germany in the Department of Philosophy, working on emotions and metaphor from a phenomenological point of view, with a scholarship of the German Research Council. Now he is a computer scientist at the ZGDV, with emphasis on integration of computer sciences and humanistic sciences.

Ian Johnson is a Senior Research Fellow at the University of Sydney, and Director of the Archaeological Computing Laboratory in the Spatial Science Innovation Unit. His work focuses on the application of GIS to archaeological and historical data, Internet databases, web mapping, the development of time-enabled mapping and GIS education for the Humanities. He joined the University of Sydney in 1990 to establish the Archaeological Computing Laboratory and since 1998 has overseen development of core technologies for the Electronic Cultural Atlas Initiative, notably the TimeMap methodology and software.

Bevan Koopman completed a honours degree in Information Technology at theUniversity of Queensland in 2002. His thesis involved the development of software tools for Indigenous knowledge management. Currently he is a Research Scientist for the Distributed Systems Technology Centre and expects to undertake an internship with the Smithsonian National Museum of the American Indian, Washington DC in 2003.

Sophie Lissonnet trained in Information Management at RMIT (Melbourne). In 2002 she was awarded an Australian Post-Graduate Scholarship to work on the Quinkan Culture-Matchbox Project as a research assistant while studying for her Masters (Research) in Indigenous Studies at JCU.

Andreas Lorenz completed his master's degree in Computer Science at the University of Kaiserslautern (Germany) in 2001, and joined the research group "Information in Context" at the Fraunhofer Institut for Applied Information Technology in SanktAugustin (Germany) in spring 2002. He commenced his PhD in the research field of multi-agent systems. His further research interests include user modeling, evolutionary algorithms, and software engineering.

Paul F. Marty is Assistant Professor in the School of Information Studies at Florida State University. He has a background in classics and computer science engineering, and his Ph.D. is from the Graduate School of Library and Information Science at the University of Illinois at Urbana-Champaign. From 1996 to 2002, he was Director of Information Technology at the University of Illinois' Spurlock Museum. His research and teaching interests include museum informatics, computer-supported cooperative work, and usability engineering. He has published several articles on various aspects of museum informatics, and teaches regular workshops on usability engineering for museum professionals.

Slavko Milekic holds a medical degree (Belgrade School of Medicine) as well as a Ph.D. in Experimental Psychology (University of Connecticut). His research interests include digital design, the psychology of human/computer interaction and building of WWW-based tools for knowledge transfer. He is the co-author and principal interface designer for the "Theory of Language", a CD ROM-based text published by the MIT Press. Slavko Milekic has received several grants from the Lemelson National Program in Innovation, Invention and Creativity for teaching academic courses in the area of psychology of human/computer interaction. He currently holds the position of an associate professor of Cognitive Science & Digital Design at the University of the Arts.

Liddy Nevile has been interested in museums and technology since establishing a classroom in the Museum of Victoria late last century (in the 80's). New forms of interactive computationally-supported environments for learning and cultural activity are of particular interest. Liddy taught law for 15 years, experimenting with computers in this context; 'experimented' with young children and computers in the early 80's; explored opportunities at schools in which every child had a personal computer in the late 80's, and was Director of the Sunrise Research Laboratory at RMIT University for most of the 90's. Current particular interests are in metadata and the accessibity of Web content.

Dianna Newman is Associate Professor at the University at Albany/SUNY and Director of The Evaluation Consortium at Albany. She has served on the Board of Directors for the American Evaluation Association; assisted in writing the Guiding Principles for Evaluators which serve as the professional guidelines for practice, and is currently on the national Joint Committee for Standards in Evaluation. She has served as evaluator for several federal and state funded technology-based curriculum integration grants and is currently developing an innovative model of evaluation that will document systems change resulting from technology-based curriculum integration in K-12 and higher education settings.

Paolo Paolini is Full Professor of Computer Graphics and Multimedia at Politecnico di Milano Italy. He is also director of HOC-Hypermedia Open Center at Politecnico. He has a Master and PhD in Computer Science from the University of California at Los Angeles (UCLA). He has been active in the following research fields: data base modelling and systems, programming languages, distributed data bases, data bases views, hypertext and multimedia models, multimedia authoring systems, multimedia application development tools. He is associate editor of TOIS-ACM. He has authored more than 70 scientific papers, and organized major scientific events (including ICHIM 2001, with nearly 400 participants from 32 countries of 5 continents). He is the chief designer of the SEE project, in charge of technical and interaction requirements.

Massimiliano Pigozzi received his honours degree in electronic engineering from the University of Bologna in 2001. From the Music Academy (Conservatorio F. Venezze Rovigo), he received a degree in the classical piano in 1994, and a degree in jazz composing in 1998. He is an IEEE member and is currently working as software designer on the MUSE project, at Ducati Sistemi S.p.A. He is the creator of the MUSE XP interface implementation framework.

Dan Porter is a Web developer and art historian. He was assistant curator at Arken Museum for Moderne Kunst, Denmark in 1996-97, before returning to England to take up a scholarship at the Courtauld Institute of Art. He joined the interpretation team at Tate Modern in London for the period prior to the opening in May 2000, and has since contributed to a number of Tate projects on a freelance basis. In 2002 he produced the i-Map project for Tate which recently won a BAFTA award. Since 1998, Dan has also been employed as a creative consultant for management consultants Cap Gemini Ernst and Young.

Nancy Proctor is the New Product Development Manager for Antenna Audio. Nancy joined Antenna Audio in 2000 with a background in designing digital publishing solutions for the arts as a founder and director of New Art online gallery and TheGalleryChannel.com She has also worked as a curator and critic, and holds an MA and PhD in art history from Leeds University. Currently Nancy's work is focused on wireless interactive guide systems, technologies to provide disabled access, and audio-visual solutions for the Web.

Anshuman Razdan is the Director of PRISM: Partnership for Research In Stereo Modeling at Arizona State University. His research interests include Computer-aided geometric design (CAGD) and Computer Graphics; NURB curves and surfaces Approximation, and use of high bandwidth networking for scientific visualization and feature segmentation. He is a PI on several NSF grants including a recent KDI grant on 3D Knowledge: Acquisition, Representation and Analysis (3DK). He has a BS and MS in Mechanical Engineering and Ph D in Computer Science. PRISM was established in spring 1996 at ASU to promote interdisciplinary research in the areas of 3D Data Acquisition, Visualization & Modeling, and Form Realization. PRISM researchers come from diverse backgrounds such as Biology, Fine Arts, Archaeology, Anthropology, Computer Science, Mechanical and Industrial Engineering, to name a few. The common interest among various disciplines of visualizing 3D and higher dimensional data is the key motivation for PRISM.

Jeremy Rowe is the Head of Media Development with Information Technology at Arizona State University. His responsibilities include involvement in institution and system level strategic planning for technology and distance education, and development of related policy including copyright, licensing, and multimedia materials ownership. Among his research interests are digital libraries and digital government. He is currently a principal investigator for NEH Moving Waters: The Colorado River and the West, and a 3-year NSF grant to develop visual query tools and interface to create a digital library for 3 dimensional research data.

Tullio Salmon Cinotti is Associate Professor at the school of engineering of the University of Bologna where he teaches fundamentals of and Microprocessor based design and Computer Architecture. His research interests are in the area of pervasive computing, mobile computing and cultural heritage based education. He is a member of ARCES, the research center on electronic systems originated by the Dipartimento di Elettronica, Informatica e Sistemistica of the University of Bologna. In 1982 he was a cofounder of Boconsult, an SME focused on system design, recently merged in Ducati Sistemi S.p.A. He is currently responsible for the Muse Project.

Sebastian Sauer is founder and head of the user interface company "ion2s". He studied Media System Design at the University of Applied Sciences at Darmstadt. During his activities he worked, authored and specialized in the field of digital experience design, information design and dynamic user interfaces for mobile information appliances. He is responsible for the underlying principles of the Transforming User Interface, Kids Innovation and momuna.

David Schaller is the founding partner of Eduweb. He is on a perpetual quest for the sweet spot where learning theory, Web technology, and fun meet. He has developed online learning activities for Colonial Williamsburg, the Discovery Channel, the Minneapolis Institute of Arts, the DuPage Children's Museum and Chicago's Brookfield Zoo, among others. His projects have won numerous awards, including two Best of the Web awards at the Museums & the Web conference and three MUSE awards at AAM 2002. Previously, as a geographer and freelance writer, he wrote museum exhibits about history, nature, and Jefferson Davis.

Kim Sheahan wears two hats at the Spurlock Museum, University of Illinois. As Special Events Coordinator, she serves as organizer and facilitator for a wide variety of Museum-based events, including lectures, panel discussions, receptions, conferences, and symposia, as well as performances by storytellers, dancers, musicians, and singers from around the world. As Assistant Director of Education, she creates, reserves, and presents outreach and gallery programs for regional community members of all ages and backgrounds. Kim is also a professional storyteller and uses stories to expand and enhance her culture-based and natural history-based programs. Kim has a B.A. in History and Classical Civilizations and an M.A. in History from the University of Illinois. She has also received an M.A. in Interdisciplinary Studies (Museology, Education, and Special Education) from the University of Idaho.

Jane Sledge has been involved in thinking about and managing museum information since 1977. She has worked with the Canadian Heritage Information Network, the Smithsonian Institution, the UNESCO-ICOM Museum Information Center in Paris, the Getty Information Institute, and is now with the Smithsonian's National Museum of the American Indian. As NMAI's Information Resources Manager, she helps to organize and connect NMAI's intellectual and informational resources and databases to users and audiences world-wide.

Oliviero Stock works at ITC-irst since 1988 and was the director from 1997 to 2002. His activity is mainly in artificial intelligence, natural language processing, intelligent interfaces, human-computer interaction, cognitive science, knowledge-based systems. He has played a key role in the launch of many ITC-irst projects funded by European research programs.

Chris Tellis founded Antenna Audio in 1986 and is responsible for developing the initial hardware and production techniques, which established Antenna Audio as the leader in the field of museum and historic site audio tours.

Shigeki Yokoi is a Professor of Graduate School of Human Informatics at Nagoya University, Japan. He received his Doctoral degree in Electrical Engineering from Nagoya University in 1977. He has been an associate professor at Mie University and an associate professor at the School of Engineering at Nagoya University. His research interests include computer graphics, virtual reality, interactive media and social affection of digital technologies.

Massimo Zancanaro received a "laurea" degree in Computer Science from the University of Milano in 1992. He joined the Cognitive and Communication Technology Division at ITC-irst in 1993. His primary research interest is in the area of Intelligent Interfaces, in particular language-based interfaces, Recently, he worked on the design issues related to mobile audio-based hypermedia. He was actively involved in many European funded-projects in the 3rd, 4th and 5th Framework.

Andreas Zimmermann completed his master's degree in Computer Science at the University of Kaiserslautern (Germany) in 2000. After one year of business work he joined the research group "Information in Context" at the Fraunhofer Institut for Applied Information Technology in Sankt Augustin (Germany) to acquire his doctoral degree. His research interests cover areas like artificial intelligence, user modeling, personalization and nomadic systems.

About the CD-ROM

This book is accompanied by a CD-ROM containing electronic versions of all these papers and many others, presented in HyperText Mark-Up Language (HTML), the technical language of the Web. It also includes abstracts of all the papers, demonstrations and workshops presented at the conference, and the biographies of all speakers and presenters. These electronic versions include color illustrations and links to the sites discussed and referenced.

You don't have to be connected to the Internet to read the papers on the CD-ROM or to navigate the full background information about the conference. You will need your own connection to the Internet to go to the linked museum sites and to follow the external links in the papers.

To use the CD-ROM, you will need a Web browser (Netscape 4.0 or Internet Explorer 4.0 or higher are recommended). Put the CD in your computer, launch your browser, and, using the File / Open menu choices, navigate to the index.html file in the main directory of the CD-ROM. Open this file in your browser – all other files are linked from there.

- **Speakers** provides a list of all the speakers at the conference and links to their abstracts, biographies and papers (where available).

- **Sessions** provides an overview of the Museums and the Web 2003 conference program and links to abstracts and paper biographies. Paper Titles link to Abstracts.

- **Abstracts** describe all MW2002 program element, including demonstrations and workshops. ***Paper Titles link to full papers when available.***

- **Best of the Web** will take you to the results of the Best of the Web 2003 conference, but you have to be connected to the Internet to do this.

If you have any questions or problems using the CD-ROM please email info@archimuse.com, and we'll do our best to help you.