NEW INTERACTION TECHNIQUES FOR THE DIGITAL LIBRARY

3D Focus+Context Interactive Visualization

A dissertation submitted by

Horn-yeu Shiaw

In partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Computer Science

TUFTS UNIVERSITY

August 2003

Adviser

Dr. Robert J. K. Jacob

ABSTRACT

In this thesis, we present a new set of interaction techniques for the digital library as a solution to the focus versus context problem where a small focus (a page of information) competes for limited screen space with a large context (the collection of information). Our solution provides for 3D focus+context (F+C) interactive visualization in virtual reality (VR). We have designed information-rich 3D worlds enhanced by lightweight interactions where users can navigate and interact with a specific datum (focus) in relation to the complete data set (context). To demonstrate feasibility and functionality, we have built three high-fidelity prototypes in desktopVR, expandable to portable immersive VR, populated with 3D data artifacts extrapolated from actual image collections in the Perseus Digital Library at Tufts University (Perseus). The first prototype is an interactive London walkthrough constructed from an overlay of several building, plot and street maps. The second prototype is an interactive Greek coin catalog in the orientation of a 3D scatter plot. The third prototype is an interactive Greek vase museum where a bird's-eye view shows a 2D scatter plot of the vase placement in the room. To evaluate our approach and measure the benefits of our techniques, we compared the vase museum with its source collection in Perseus. The results of the experiment showed that users who used our prototype learned more about the same set of information nearly three times faster and with significant improvement in accuracy.

ACKNOWLEDGMENTS

I would like to take this opportunity to thank all my teachers, committee members, peers, friends and family for their help and support in making my study possible.

My deepest gratitude goes to my adviser, Prof. Robert J. K. Jacob, for his inspiration and tutelage. He has instilled in me a profound interest in the subjects of my research and a mental model of what a good scholar should be. I harbor a humble respect for my adviser whose footsteps I hope to follow.

I am grateful to Prof. Anselm Blumer for his help in making my study achievable. I am grateful for his help in the summer of 2001 in tutoring me for the algorithm and theory qualifying exams.

Special thanks to the entire body of Tufts' Perseus group, headed by Dr. Gregory R. Crane of the Classics Department, where I have been a research assistant since 2000. The Perseus Digital Library project is sponsored by the National Science Foundation and the National Endowment for the Humanities, Digital Libraries Initiative Phase 2 Program.

I am grateful to Prof. Steve A. Morrison. He is my academic predecessor, a committee member and the professor of two courses for which I was a teaching assistant (Comp 180: Software Engineering in Fall, 1999 and Comp 15: Data Structure in Spring 2000).

I thank Dr. Martin Wattenberg, a committee member and a researcher at IBM, for his expertise and recommendations. Dr. Wattenberg's research interests include information visualization and its application to collaborative computing, journalism, bioinformatics, and art.

I would like to thank my peers at Tufts University, who helped pilot and test the prototypes in various stages of development. Their valuable input had a significant effect in the design and implementation of the final version.

I thank the 20 anonymous test subjects who offered their help in the usability test of the vase museum.

Our research is supported by the NSF Interactive Systems Program, Office of Naval Research, Berger Family Fund and the Tufts Selective Excellence Fund. We greatly appreciate their support.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

New Interaction Techniques for the Digital Library

3D Focus+Context Interactive Visualization

CHAPTER 1

THESIS INTRODUCTION

The worldwide web (WWW) is an effective provider of digital data. Unfortunately, it is also a notorious hider of contextual information. There is a constant trade-off between '*a view of the whole data available, while pursuing detailed analysis of a part of it*.' [164, pp 307]. This is known as the **focus versus context** problem where the overview and detail compete for limited screen display. One solution to the focus versus context problem is referred to as the **focus+context** (F+C) technique where the **user interface** (UI) design affords a strategic balance between the local (focus) and the global (contextual) information. According to S. K. Card et. al. [164, pp 307], there are three premises that can define a solution as focus+context in nature: '*First, the user needs both overview (context) and detail information (focus) simultaneously. Second, information needed in the overview may be different than that needed in detail. Third, these two types of information can be combined with a single (dynamic) display, much as in human vision.*'

In other words, a solution that is focus+context in nature must be able to demonstrate the three premises. First, there is a *need* for both focus and context simultaneously. Second, focused information is *different* than contextual information. Third, focus and context are provided for in a *single display*. In this thesis, we highlight the focus versus context problem specifically for the web-based **digital library** (DL). Many digital library collections offer vast amounts of information. The users need to see both the focus (a detail document) and the context (the collection that the detail document belongs in) simultaneously. The information represented in the detail document (focus) is generally different than that of the collection (context). On the web, the users face the problem of viewing an isolated page (focus) of information while endlessly branching out to the cyber links (more context). A document window (one page of information) can take up the bulk of the screen space while the contextual window (generally a meaningful list of documents) is relegated to the background. The users can easily lose focus or context. Therefore, there is a need to show both focus and context in a single display.

In this thesis, we present a solution to the focus versus context problem for the digital library. We provide for a new set of **non-WIMP** (Window, Icon, Mouse, Pointer) [76, 96] interaction techniques in 3D focus+context interactive visualization. Our methodology deviates from the traditional approaches in better **GUI** (Graphical User Interface) design, more compact 2D screen layouts, greater retrieval efficiency, advanced browsing techniques or hierarchical semantics. [1, 2, 4, 6, 24, 40, 48] Instead, we provide for suggestive (tacit, non-command) techniques in **virtual reality** (VR) [44] where the users can interact with data objects in an information-rich environment. The 3D models used to populate the virtual environments (VEs) are constructed from large data sets found in a real digital library. [73, 74, 76, 81, 85, 99] Specifically, we are creating a 3D information landscape where the users can focus on a piece of datum by navigating in the virtual scene while the context of the entire data available is maintained in the surrounding environment. We have designed the 3D world based on the contextual nature of the information and populated the landscape with realistic 3D models of the data artifacts. We allow the users to explore dynamically in our virtual environments, as we display additional layers of information based on **lightweight interactions**. Our solution differs from the traditional task of reading and taps into the visual ability of the average users. We strive for a duplex communication between the system and the users in a virtual landscape.

To validate our design concepts, we have implemented three high-fidelity prototypes in virtual reality. Our prototypes are functional in desktopVR and expandable in portable immersiveVR. The 3D models used to populate the virtual scenes in the prototypes are constructed from actual collections found in the **Perseus Digital Library** (Perseus) at Tufts University [52, 63]. The first version of Perseus was a CD-ROM containing a set of hypertext pages. Today, Perseus has become a leading web-based digital library in classics and Greek archaeology. It contains many image collections of old maps, ancient sites, mythologies, gems, coins, vases and sculptures. The development, evolution and social effects of Perseus have been the topics of many research studies. [54, 55, 95, 137, 142, 143] For our prototypes, we have selected the London maps, the Greek coins and the Greek vases as our representative data sets. We have made these selections to take advantage of the complex and non-textual nature of the data content. Although we have used Perseus as our data source, the solution proposed in this thesis is applicable to digital library in general. The main page of the Perseus web site is shown in FIGURE 1.1. From the main page, the users can access the London Bolles Collection [30] of the old London maps and the Greek Archaeology of the coins and vases used in our prototypes.

FIGURE 1.1 PERSEUS MAIN PAGE http://www.perseus.tufts.edu/

The three prototypes we have built are: 1. the **3D interactive London walkthrough** (London walkthrough) constructed from a set of old London maps in the Perseus archive; 2. the **3D interactive Greek coin catalog** (coin catalog) modeled from two Greek coin collections in Perseus; and 3. the **3D interactive Greek vase museum** (vase museum) selectively modeled from the London vase collection in Perseus. Collectively, these three prototypes intend to show that focus and context co-exist in a single display. Large data sets can be visualized as an information landscape (context) where the users can navigate and 'walk to' a datum of interest (focus).

The London walkthrough is a 3D cityscape of 1840's London (Section 3.3.1, Chapter 4, Appendix B.1) where the users can navigate in the virtual environment interactively. FIGURE 1.2 shows a rendition of the London walkthrough in an eye-level view.

FIGURE 1.2 EYE-LEVEL VIEW OF THE LONDON WALKTHROUGH

The virtual environment is reconstructed from a set of old building, plot and street maps. There are 242 building-like boxes, texture-mapped with images of the buildings cut out from the original 2D building maps (Appendix B.1). The buildings are lined up to form an intersection. The intersection is bent in 3D space to match the contour of the street map. We are able to combine three types of maps (i.e. building, plot and street maps) in a new design technique called the **3D map overlay** (Section 3.4.1, Section 4.2). FIGURE 1.3 shows a rendition of the London walkthrough in a bird's-eye view, where the roofs of the buildings are oriented in a way that can reflect the plot map over the street map.

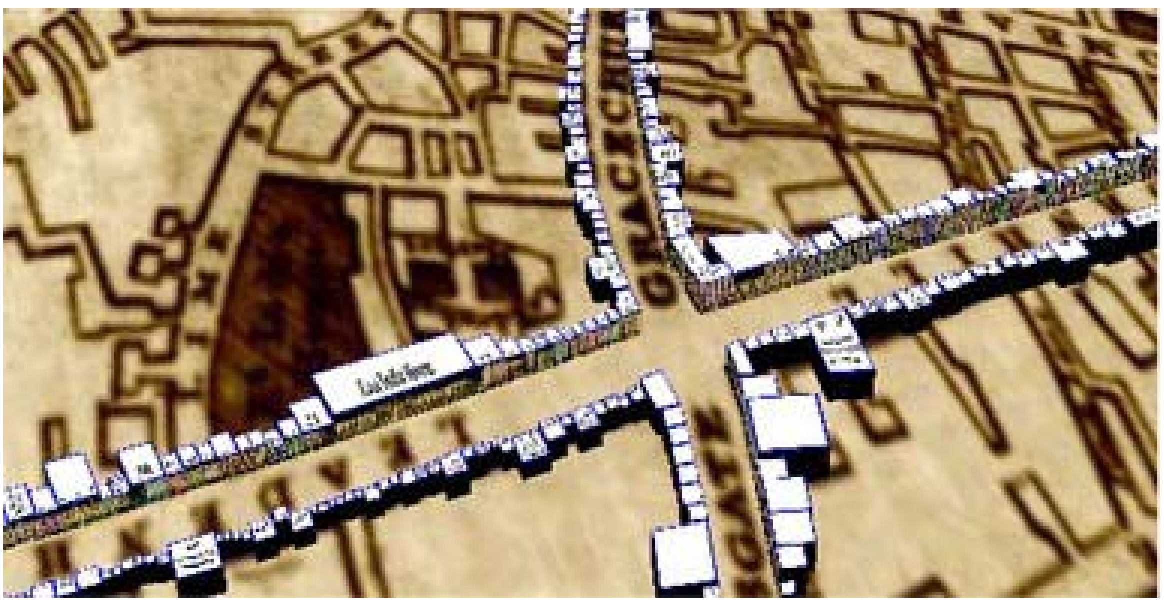

FIGURE 1.3 PERSPECTIVE VIEW OF THE LONDON WALKTHROUGH

In a single virtual display, we are able to provide for the context (the cityscape) and the focus (a building) by navigation in XYZ. Using another rendition of the London walkthrough, we are able to show a London fire that occurred in that intersection in 1765.

The users can explore the burnt scene, as compared to reading about the damage in a text passage in Perseus (Section 5.4, Appendix B.2). We can also allow the users to interact with our **3D text menu**, which is analogous to hyperlinks on the web, where the users can read 3D floating text and select an embedded word as a menu option. The system will then 'take' the users to the relevant location in a virtual scene linked to the word selected.

Our second prototype is the coin catalog. The coin catalog is a 3D room populated with two collections of 132 Greek coins found in Perseus (Chapter 5, Appendix B.3). One collection contains 62 **Dewing** (Dewing Numismatic Foundation) coins. The other collection contains 70 **BCMA** (Bowdoin College Museum of Art) coins FIGURE 1.4 shows a rendition of the coin catalog in an eye-level view. The right cluster is the Dewing coins and the left cluster is the BCMA coins. The 3D coin models are animated in automatic rotation.

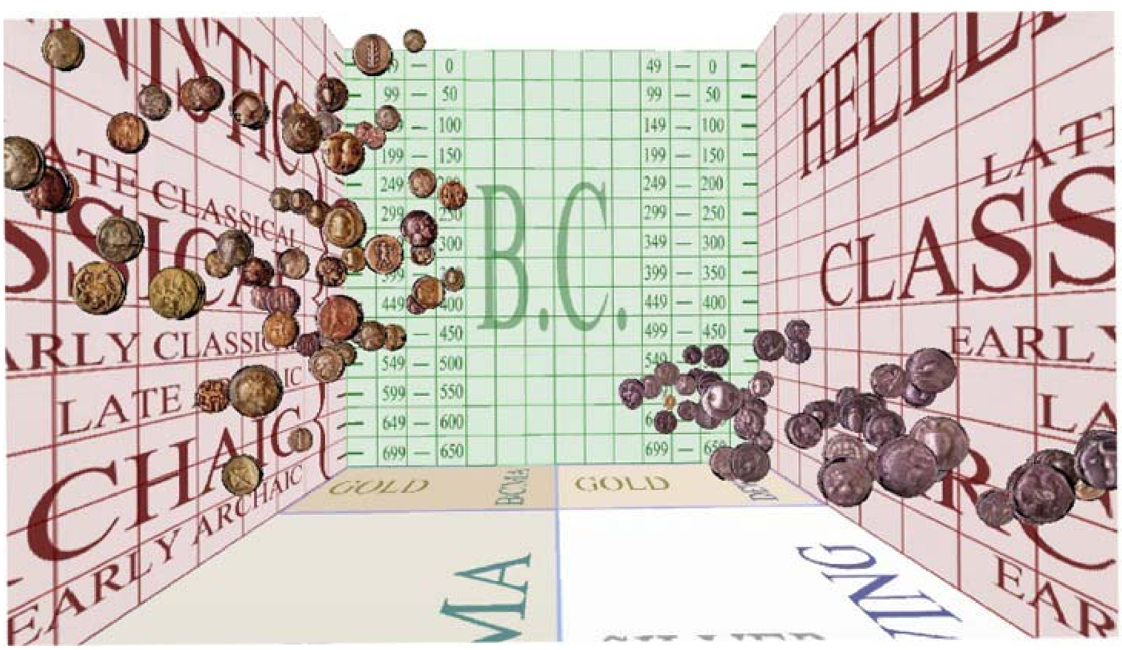

FIGURE 1.4 EYE-LEVEL VIEW OF THE COIN CATALOG

The 3D room has the layout of a XYZ coordinate system. The X- and Y-axes form the floor of the 3D room. One side of the floor (X-axis) is divided into two regions, one for each of the two collections. The other side of the floor (Y-axis) is divided into three regions, one for each of the three metals: gold, silver and electrum (an alloy of gold and silver). The height of the room (Z-axis) shows the year in B.C. and the corresponding historical periods (e.g. archaic, classical). A coin is placed in the 3D room according to its collection, material and year/period. The placements of the coins in the 3D room form a **3D scatter plot** (Section 3.4.2, Chapter 5.2). FIGURE 1.5 shows another rendition of the coin catalog in 3D perspective. In this rendition, the grid pattern is removed and the background is grey to reduce color noise.

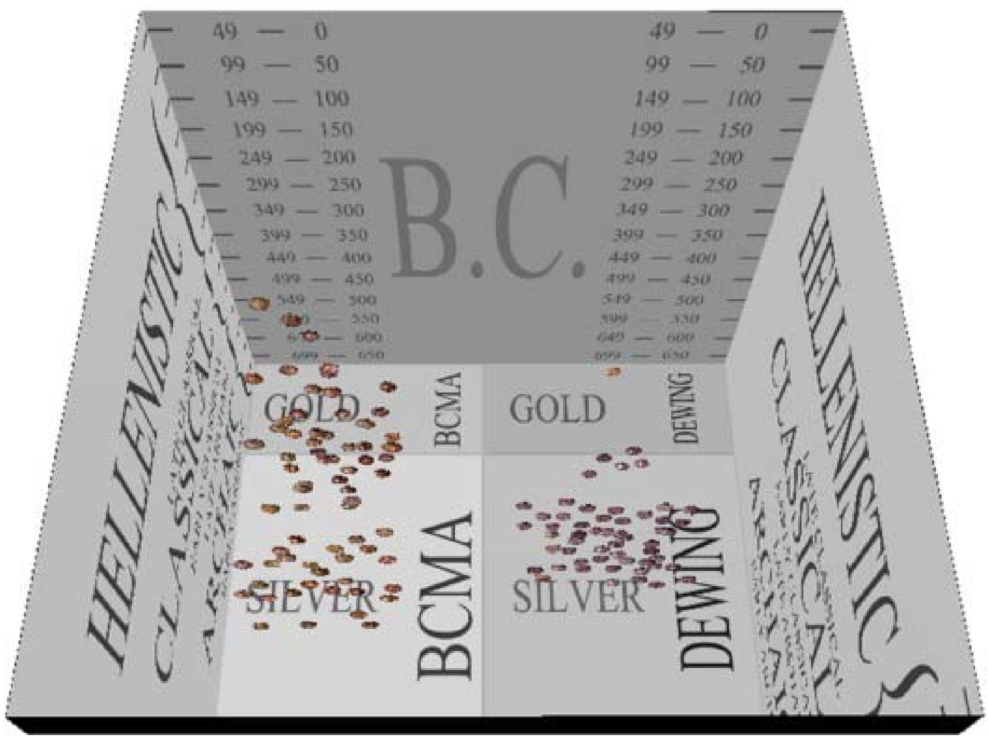

FIGURE 1.5 PERSPECTIVE VIEW OF THE COIN CATALOG

From a bird's-eye view, the floor of the 3D room and the placement of the coins show another 2D graph by collection and material. The coin catalog allows easy comparison of the two collections. As the users navigate in the coin catalog, both focus (single coin) and context (two collections) are maintained in a single display. In Perseus, the current approach is to browse and read the corresponding web sites, which are listed by the catalog numbers of the coins (Appendix B.3).

Our third prototype is the vase museum. The vase museum is also a 3D room, which displays 157 Greek vases (Section 3.3.2, Chapter 6, Appendix B.4). These vases are selected from the London collection in Perseus. One side of the wall (X-axis) shows the production year in B.C. The other side of the wall (Y-axis) shows the wares (types of glazing; e.g. red figure, black figure). FIGURE 1.6 shows an eye-level view of the vase museum.

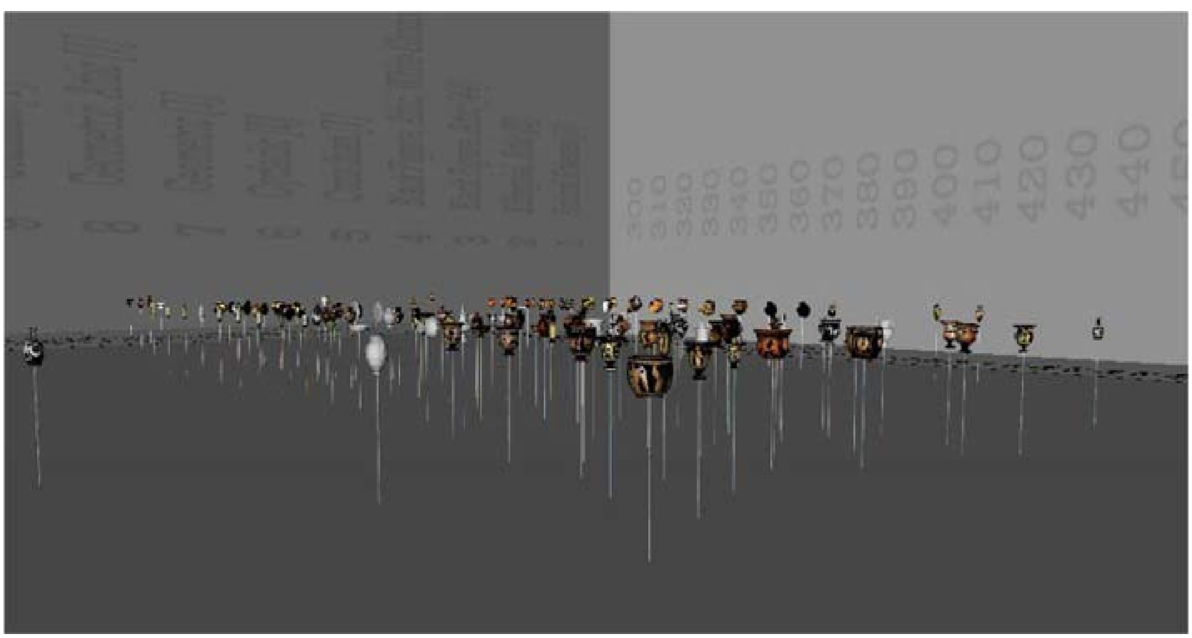

FIGURE 1.6 EYE-LEVEL VIEW OF THE VASE MUSEUM

From a bird's-eye view, the vase museum shows a 2D scatter plot by years and wares. FIGURE 1.7 shows a perspective view of the vase museum in a bird's-eye view.

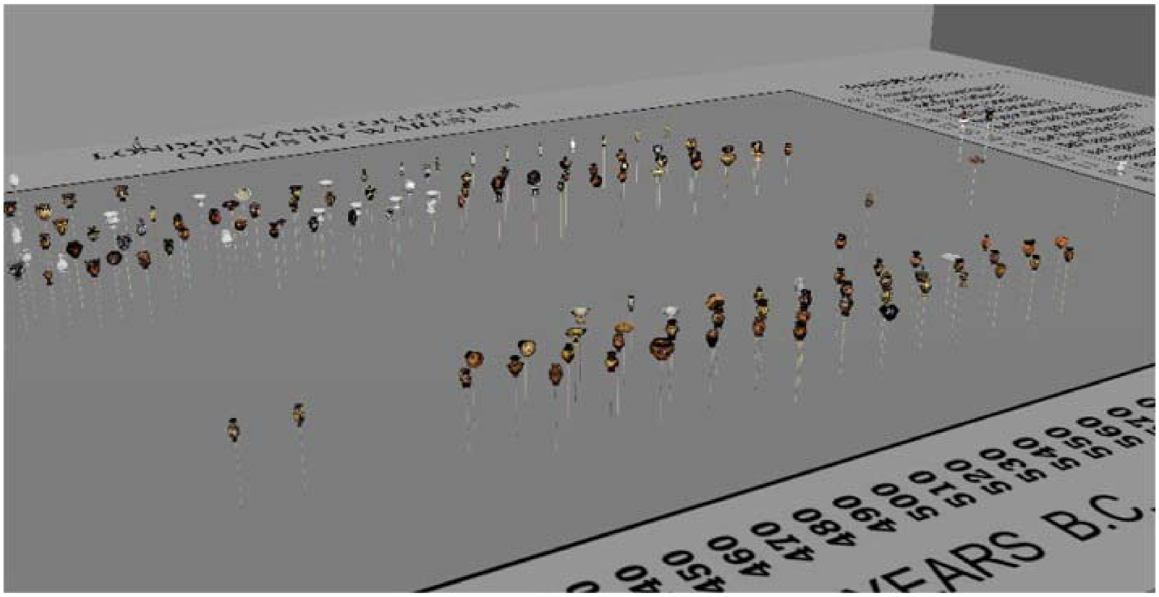

FIGURE 1.7 PERSPECTIVE VIEW OF THE VASE MUSEUM

As the users navigate in the vase museum, the context of the collection is maintained in the background while the focus is placed on a single vase. The users can compare a vase with its neighbors or vases in the background. Using preset camera views, the users can compare two groups of vases. The 3D models of the vases are modeled from available 2D photographic images in Perseus. Proximity to a vase triggers its supporting information (e.g. captions) to be displayed in the virtual scene. When the supporting information is displayed, the vase museum allows users to switch to the 2D HTML (HyperText Markup Language) pages while maintaining their position in 3D space. Animated tour, called **anitour**, can take the users on a guided tour in the museum.

In an IRB[1]-approved usability test, we are able to evaluate the benefits of the vase museum by comparing it to the source site in Perseus (Chapter 7, Appendices C and D). We have two different groups of 10 users each (between subjects). One group used the Perseus digital library (searched and browsed on the web) and the other group used the vase museum (navigated in a virtual environment). Users in both groups were given the same list of 10 vase-related questions while they were using the system. The amount of time used to answer each question was recorded. These 10 questions are designed based on real homework one would find in archaeology courses. The nature of the questions reflects the overall features of the collection (e.g. number of handles, general appearance, colors, etc.) and requires the users to focus on a vase while comparing it to the other vases in the collection (context). Both groups of users took the same quiz at the end of the test. The results of the experiment showed that the group who used the vase museum is able to perform the same task with improved accuracy in nearly ⅓ the time.

1.1 THESIS OVERVIEW

The chapters in this thesis are organized as follows:

Chapter 2 will review related work in human-computer interaction techniques for the digital library. We will examine current issues in displaying large data sets in the digital library. Existing literature in focus+context will be discussed and compared. We will

[1] Institutional Review Board for research involving human subjects. 'IRB 101 on the Road" can be found at: http://www.primr.org/101road.html

review popular metaphors, such as trees, maps and books, which are frequently used to represent a large context in the digital library. We will look into zooming techniques that can bring focus on a large context by distorting a target area. Since our solution is 3D-based, we will also examine existing approaches in virtual reality, including research in 3D and 3D-like navigation and exploration.

Chapter 3 will propose solution concepts to resolve some of the issues in the focus versus context problem for the digital library. Theses issues are cluttering, competition for screen space, deficiencies in the 2D paradigms, branching and connectivity confusion. The design and implementation of our solution methodology call for 3D focus+context interactive visualization in either desktopVR or immersiveVR using lightweight interactions. Our proposed design concepts include adhering to the true nature principle, presenting information art (InfoArt), creating 3D interactive walkthrough, 3D interactive museum, anitour, 3D map overlays, 3D interactive graph (e.g. 2D and 3D scatter plot), and 3D text menu.

Chapter 4 will present the design rationale and implementation issues of the 3D interactive London walkthrough prototype. The London walkthrough is a 3D virtual rendering of an intersection of 1840's London and the fire that occurred there in 1765. The original source maps used can be found in Appendix B.1. Implementation issues such as sizes, scales, level of details and loading concerns will be discussed. We will compare the current approach (reading about the fire) and the new approach (exploring the 3D walkthrough of the fire) in an application of the London Fire of 1765. A text description of the London fire can be found in Appendix B.2.

Chapter 5 will present the design rationale and implementation issues of the 3D interactive Greek coin catalog. The coin catalog is a 3D scatter plot populated by 3D coin models reconstructed from two collections of Greek coins in Perseus. We will compare the current approach of browsing the HTML pages of the coins with our new approach of visualizing the 3D coins in a virtual room showing two clusters of collections. The two collections of Greek coins modeled can be found in Appendix B.3, listed by their catalog numbers.

Chapter 6 will present the design rationale and implementation issues of the 3D interactive Greek vase museum. The vase museum has on display a collection of 3D Greek vases found in Perseus. The current approach in Perseus shows the collection as a list of catalog numbers from which the users can branch out to corresponding vase HTML pages. The users learn about the vases by reading the descriptions and viewing available pictures. The new approach shows all vases in a 3D room-like museum. The users can visit the 3D museum and learn about the vases with little or no reading. The complete list of the catalog number of the Greek vases modeled can be found in Appendix B.4.

Chapter 7 will present the usability test and the evaluation of the vase museum. The testing metrics, strategy, issues, results and summaries will be carefully examined. The complete set of data collection instruments (forms, surveys and questionnaires) can be found in Appendix C. The actual data collected are summarized in Appendix D. Our test results show that our new interaction techniques are better than conventional ones for performing a focus+context task for the digital library.

In the summary sections of Chapter 4, 5 and 6, we also present possible future directions in research. Some of the major areas for future improvements are faster rendering speed, advanced modeling algorithms, full immersion, database management of 3D models, and more interaction techniques. Finally, Chapter 8 concludes the thesis.

The designs of the three prototypes require realistic data artifacts to be constructed true-to-nature and artistically. The prototypes call for 242 3D buildings, 132 3D coins and 157 3D vases to be modeled and placed in special virtual layouts. In order to create the virtual scenes in a reasonable amount of time, we have devised component prototyping techniques to rapidly model and render complex scenes in desktopVR. We have also shown feasibility of portable immersiveVR on a laptop. The appendices will present the prototyping environment (Appendix A), source data (Appendix B), data collection instruments (Appendix C) and data collected (Appendix D).

For the rest of this chapter, we will give a brief background review and present our research motivation, problem refinement and contribution.

1.2 BACKGROUND

The work presented in this thesis is multi-disciplinary in human-computer interaction (HCI), information visualization (InfoVis), virtual reality (VR) and the digital library (DL).

Specifically, we are interested inventing non-WIMP and lightweight interactions in virtual reality where the users can focus without losing context. We seek to invent new interaction techniques that can benefit from new ways of visualizing both concrete and abstract data in the digital library. We seek to merge the benefits of the different disciplines required for our proposed solution. FIGURE 1.8 shows the discipline diagram for our design concepts.

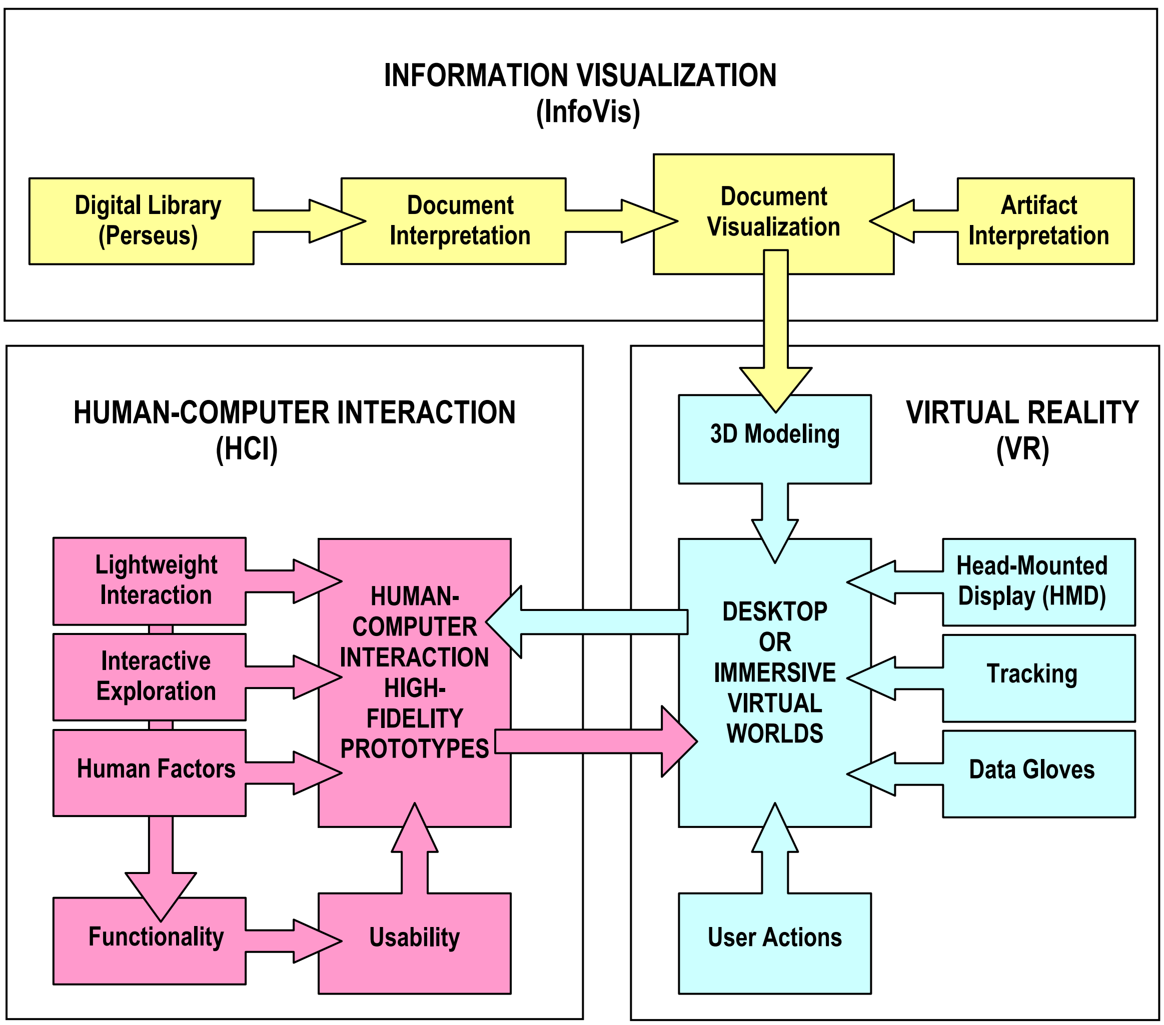

FIGURE 1.8 DISCIPLINE DIAGRAM

Human-computer interaction is a relatively new field that incorporates human factors in the design of a more functional system. [10, 133, 158, 163, 167, 181] The Curriculum Development Group of the **Association for Computing Machinery** (ACM) and **Special Interest Group on Computer-Human Interaction** (SIGCHI) defines human-computer interaction to be '*a discipline concerned with the design, evaluation and implementation of interactive computing systems for human use and with the study of major phenomena surrounding.'* Human-computer interaction was not a major concern in the early days of computing machines. The lack of interest in human-computer interaction was due to the low expectations in usability and the high skill levels of target users. The systems are meant to be used by engineers, scientists, special interest groups and computer professionals who are willing to invest the time and energy to learn a cumbersome system. It was not until the late 1970's that the interfaces of computer systems shifted from engineering-centered to user-centered. Today, it is well recognized that substandard user-interface design can be a major contributor to user dissatisfaction, ineffectiveness and human errors. As part of the human-computer interaction design process, high-fidelity prototypes can be built before committing to a full-fledged implementation. Usability testing is frequently done in different stages of development to ensure the soundness of the design and implementation. [157]

Information visualization is also a relatively new field of research established in the early 1990's with conceptual inception dating backing much earlier. Today, information visualization has far-reaching effects in a large number of subject domains, including the digital library and human-computer interaction. There are many different types of information visualization, such as scientific, geographical, astronomical, statistical, abstract, and document visualization. Even though human perception of images is not fully understood [119], the fundamental concept in information visualization lies in the hope that 'a *picture is worth a thousand words*' [125] and the goal of '*using vision to think.*' [164] Cognitive psychology and human vision also play important roles in the definition of effective visualization and affordable interactions. [145, 178, 182, 194, 195, 205] The effectiveness of a given visualization depends on issues such as the selection of raw data, the accuracy of interpretation, the correctness in implementation and the power of the delivery techniques. Good visualization techniques should facilitate the process of learning and retention. However, learning can be harder to measure than retention. Many visualization techniques are symbolic, geometric, or metaphoric. Some symbols, tables or diagrams are just as hard, if not harder, to learn than any written languages. The digital world offers a data-rich environment that is no longer confined by text and pictures. Some of the topics in information visualization are cross-referenced in human-computer interaction probably because effective interaction techniques call for effective visualization and vice versa. Current focus+context solution techniques tend to be graphically visual (Chapter 2) in implementation because interface designers are exploring techniques beyond reading text on a 2D screen. Designers are looking into the intrinsic nature of information representation.

Virtual reality has its origin in flight simulations dating as far back as the 1950's. Virtual reality remained dormant until the mid-1980's when it hit the headlines after real-time computer graphics became commercially affordable. At the time, only **immersive**

virtual environment (IVE) or immersiveVR was considered authentic virtual reality. Major components of a virtual reality system may include a **head-mounted display** (HMD) [192], tracking devices, data gloves and tactile sensors. Eye-scanning with or without HMD can be used to calibrate the direction of gaze. [97, 98, 193] During the 1990's, the concept of desktopVR or **fish tank VR** was gradually accepted as a scale-down version of the IVE. Since then, virtual reality has evolved and diversified into many areas of research and development, including human-computer interaction, information visualization and the digital library. Both desktopVR and immersiveVR require faster computer speed, better graphics and improved visualization paradigms. The potential for virtual reality applications are limitless but questions remain whether a viable system could ever be built. Before the technology is fully matured, the users must have reasonable expectations. Some may believe that virtual reality technology has progressed much more slowly than what was originally expected, may not be worth the effort, or may not go far. Nonetheless, virtual reality will remain a frontier in many research areas because of its unleashed potential in immersive realism.

The concept of the digital library was first formalized in 1940's by Vannevar Bush in his **Memex** where he had the contents of all his books, records and communications on microfilm. Since there was no digital storage at the time, his concept was more futuristic than realistic. But the idea continued to grow in the 1940's and 1950's. In the 1960's, MIT developed Project **Intrex** when digital storage became available. In 1978, a book called "**Toward Paper Information Systems"** by Lancaster first described the concept of the digital library as we know it today. The digital library could not have become an academic and commercial reality without the worldwide web. The year 1994 was an important year in the history of the digital library. In that year, the Library of Congress launched the **National Digital Library** (NDL) Project. In the same year, **Digital Library Initiative Phase I** (DLI-1) was launched by the National Science Foundation (NSF), National Aeronautics and Space Administration (NASA), and the Defense Advanced Research Projects Agency (DARPA). Phase I focused on making the digital library an everyday phenomenon. Research efforts concentrated in the development of techniques in the collection, storage, retrieval and communication of the digital library. The digital library can continue to benefit from technology improvements in many fields of research, including human-computer interaction, information visualization and virtual reality. [148, 149, 160, 161, 172] In 1998, Phase II of the Digital Library Initiative (DLI-2) was launched. By this time, the digital library had become a reality. Phase II of the initiative is concerned with the usability of the digital library, which in essence places the development of next-generation interaction techniques on the formal path of research in the digital library. [112]

1.3 MOTIVATION

Our motivation to create better interaction techniques derives from the need to offer the users more effective interactions with the information buried in the digital library. Prior to the era of the web-based digital libraries, information was primarily preserved in manuscripts and printings. Physical proximity was crucial for information access. Today, the users can type in a few keywords and obtain lists of matching results. In the beginning, the simple fact that this could be done was exciting enough. Soon the initial enthusiasm was replaced by the overwhelming frustration of plowing through endless hyperlinks. As the number of documents grows exponentially, searching on the web is becoming a time-consuming task. [33] The search engine is generally able to satisfy the keyword constraints but unable to provide for any semantic relationship in the matching documents. The task of contextual interpretation is left to the users. Some systems compensate for this problem by categorizing similar materials into selectable collections to bypass the need for contextual query. Such solution is patch-up in design. Part of the cause of the problem was historical. When information was first stored digitally, the primary challenge was in establishing the infrastructure of the digital library. Organizational design was based on technological expediency and empirical solutions. Collections were organized in files and directories. Documents were simply displayed as HTML pages with or without images. There was relatively little effort in the contextual design or structural planning of information semantics.

1.4 PROBLEM REFINEMENT

In the prevailing desktop metaphor, the user-interface techniques for the digital library can be characterized as a process of search, browse and display. As the users browse the web pages, the contextual purpose becomes a moving target as they click on a hierarchy of web-links. The root of the obstacles that hinder the development of advanced and next-generation interaction techniques lies in the way information is visualized. A poor visualization can obscure the nature of the data. Current interpretation and representation of large library data sets do not permit the users to see both the trees (focus) and the forest (context) [63]. The context is generally given by a list of search results or a display of organized topics. The focus is generally in the form of a single HTML page. User-interface designers face the challenging task of balancing focus and context in a single display. The users can see either the focus or the context but not both at the same time in the same screen space. As the use of the digital library has become an important part of our life, the demand for better interaction techniques beyond the desktop metaphor is becoming more pressing.

A major hurdle to overcome in solving the focus versus context problem for the digital library is screen contention. Due to the limited 2D space on the screen, focus (a HTML page) and context (the collection) are competing for a very small view port. There are also the issues of occlusion and cluttering in a large context. Whereas a single document can take up the entire screen, a large context in a single display can be hard to read. In that respect, virtual reality is a promising platform because it offers more virtual screen real estate. Another great benefit of virtual reality is the sense of presence. The users can potentially be given the opportunity to see, hear, touch, feel, explore and interact with a computer-generated world whether or not the real counterpart is available. However, in virtual reality, there is the problem of translating printed text into visual representations in 3D landscapes. Virtual scenes are costly and time-consuming to build. The challenge is to examine existing HTML pages in Perseus and rearrange the information contained in them in a contextual manner. The new way of focusing in any part of the context should be intuitive and natural. Since our solution is primarily visual, the users must be able to recognize the data with little or no reading. The users must be able to maintain focus and context while navigating in a large virtual space. During the last two decades, much research effort has been directed to an information environment that can afford both focus (a document) and context (in a large data set). [25, 28, 29, 75, 92, 93, 123, 124, 162] In this thesis, we address the following research questions that have not been explored or solved to the best of our knowledge:

- Can a new set of techniques solve the focus versus context problem in the digital library?
    - Can we define a large 3D virtual context in a meaningful way?
    - Can we define 3D virtual focus on a large 3D contextual space?
    - Can the users focus on selected 3D virtual data without losing the virtual context in a single display?
- Can we implement prototypes of our proposed solution that can be run in real-time?
    - Can we implement a metaphor realistically enough that the users can interact with our 3D world with little or no training?
    - Can the 3D data in our virtual worlds be realistic enough such that the users can recognize them and learn about the information represented in them?
    - Can the users interact with 3D data in a 3D information landscape?

    - Can the 3D world arrange and rearrange to react to dynamic user interaction?
- Can we compare our implemented solution to a real-life digital library and come up with a fair assessment of the effects in both systems?
    - Can we conduct a usability test that compares both systems without bias?
    - Is our approach better given a focus+context task?
    - Can retention or learning be improved and quantified with our proposed solution?

1.5 CONTRIBUTION

As a solution to the focus versus context problem, we seek to design 3D worlds that can facilitate direct interaction between the users and the information space. The set of new interaction techniques provides for 3D focus+context interactive visualization of a large data set in a virtual reality information space, enhanced by lightweight interactions, where the users can view, navigate and interact with specific datum (focus) in relation to the complete library collection (context). We have designed and implemented virtual worlds populated with the data artifacts found in Perseus, a real-life digital library used worldwide. The design for our solution calls for a new way of visualizing and interacting with information in virtual reality. We have implemented our high-fidelity prototypes built by components. We believe that the users can learn about the same amount of information if the data can be seen in 3D focus+context, as compared to reading and viewing pictures.

We seek to improve upon the existing system by converting text and images into 3D information landscapes. The 3D models we have created can be used as value-added attributes to the current images. We also seek to expand upon the existing system by creating a methodology in focus and context that will allow the users to interact with information in a large data set. Our design and implementation seek to create a desktopVR environment that can offer immediate benefits to the existing digital library. Our desktopVR scenes can be run from third-party browsers. We also seek to prototype an immersiveVR environment that is less lab-oriented and more portable. The design concepts for the immersiveVR are not quite achievable given current computers. Although we have built high-fidelity prototypes to demonstrate our design concepts, a fully functional system is currently beyond our reach because our new set of interaction techniques puts a heavy burden on available computer resources. However, through our efforts, we hope to show a glimpse of the potentials of next-generation, interactive visualization in virtual worlds.

The contribution of the work described in this thesis can be summarized below:

- We have highlighted the focus versus context problem in the digital library.
- We have invented a set of interaction techniques to solve the focus versus context problem in several areas of our digital library.
- We have validated the feasibility of our techniques in our three high-fidelity prototypes (London walkthrough, coin catalog and vase museum).

  - Our set of new techniques offers 3D focus+context interactive visualization in virtual reality. In our prototypes, we are able to show the context of the whole data available (a set of maps, two coin collections and one vase collection) without losing focus of the detail information (a building, a coin and a vase).
  - We have demonstrated feasibility by implementing high-fidelity prototypes with functionality of our new techniques. Our prototypes are functional in desktopVR and expandable to immersiveVR.
  - We have modeled our virtual worlds from real data found in an existing digital library, specifically Perseus.
  - We have compared our new set of interaction techniques to Perseus and showed that our techniques can be used to improve upon the existing system. The London walkthrough and the 3D models of coins and vases should have immediate benefits to the current HTML displays.

- We have designed a usability test that can fairly assess the advantages of applying our new set of interaction techniques.
  - We have designed our data collection instruments analogous to real-life homework questions for archaeology courses found on the web.
  - We have successfully conducted a usability test using one of our prototypes and comparing it to its original source site in Perseus.
  - Users are able to perform a real-life digital library task using one of our prototypes. Our usability test showed that the group of users who used the virtual

vase museum is able to complete the same task given to another group that used the existing Perseus digital library.

- Our experiment showed that our new set of techniques could help the users perform the same learning task in less time with improved accuracy.
    - Our design and implementation can be used for focus+context tasks in real life. Users are able to answer questions that are contextual in nature (e.g. shape and color), without reading, simply by 'looking around' in the virtual vase museum.
    - Our prototypes can be run in real-time with minimal training and with a very short learning curve. The vase museum allows the users to carry over real-life experience. 'Looking around' in the virtual vase museum is analogous to visiting a real museum. Our group of users quickly learned what they needed to do, with little or no training.
    - Users with no prior exposure can quickly learn our system (5–10 minutes) and perform a real-life task. Our group of users has never been exposed to desktopVR or the vase museum. However, they are able to quickly learn the system and out-perform the other group of users who used a web-based digital library.
    - We observed that the users feel that they have less of a need to read and to count when exposed to a large data set.
- We have improved upon the existing paradigms and applications in our digital library.

- The techniques we have introduced supplement existing 2D photographic images with 3D virtual models.
- Our techniques in rapid component prototyping allow a large quantity of 3D models and virtual scenes to be built in relatively less time than direct programming.
- We have made a desktopVR version of our prototypes easily accessible to existing applications.
- We have made an immersiveVR version of our prototypes portable and feasible.
- We have created a design that can continue to benefit from improved computer speed and graphics power.

CHAPTER 2

RELATED WORK

In this chapter, we discuss related work in focus+context (F+C). Our original examination in related work was to compare our work to similar research in interaction techniques for the digital library (DL) in virtual reality (VR). Yet, we found our work to be new and unique for many reasons. First, we highlighted the focus versus context problem as applied *specifically* to the digital library (DL). Second, our proposed solution intersects disciplines in human-computer interaction (HCI), information visualization (InfoVis), virtual reality and the digital library. Third, our presentation of focus and context in virtual immersion is very different from the current approaches in screen displays. Fourth, our data samples are non-textual. Finally, our work is measurable and verifiable in usability.

Furthermore, the interaction techniques in our 3D virtual worlds are user-driven. We allow users to navigate freely and interact directly with objects in the virtual scene. They can focus on a specific data in the virtual scene by walking closer to it. In the meantime, they can still maintain the context of the entire information space around them. The current interactions techniques are data-driven in the sense that user actions involve visual mappings (e.g. branching and dataflow), data transformations (e.g. dynamic queries) and view transformations (e.g. direct selection and zooming). [164, pp 233] Therefore, our proposed solution has no direct single ancestral work. However, we can still survey existing literature and present related work that *could be* applied to the digital library in the context of this thesis. We have applied the following criteria in the selection of related work presented in this chapter: 1.) the context (data sets) must be significantly large; 2.) the focus (one datum) poses problems in focus versus context; 3.) the solution domain must be focus+context or must be highly similar in nature; 5.) the implementation could *conceivably* be applied to a real-life digital library project (even though it has not yet been); 4.) the visualization must be highly graphical (as compared to simple text and documents); 5.) the user interface cannot be overtly simplistic (i.e. a command-line system or a simple dialog with two buttons).

2.1 LITERATURE INTRODUCTION

The challenge of overcoming the focus versus context problem is to maintain and balance selected pieces of information with respect to overall content. This is generally achieved by choosing an effective visualization to represent the contextual information. The interactions should then allow users to bring their focus to an area of interest with minimal loss of context. Visualization styles can be categorized as geometric, metaphoric or a hybrid. Frequently used geometric interpretations are 2D or 3D trees, charts, tables and graphs. Frequently used metaphoric representations are the desktop, maps and books. Geometric visualizations are very effective in representing statistical, relational and hierarchical information. For example, the **table lens** [153] provide for a focus+context fisheye technique that is effective in visualizing tabular information. The **tilebars** [82] are effective in showing the presence of certain documents using shaded bars. User interactions generally result in the rearrangement of the geometric layouts. It is intuitive that no single visualization style is sufficient in representing all types of information in the digital library. Selecting the right visualization tool for the intended interactions is important. The goal of all visualization is to effectively convey the key features of a complex structure by powerful representations. Advancement in interaction techniques allows users to communicate with the computer in a way that can facilitate the absorption and digestion of information [31]. Ease of navigation is a research concern in both 2D and 3D screen space. [15, 70, 140, 141, 187, 197] Therefore, the interactions serve the purpose of arrangement and rearrangement of information selected and displayed. In light of that, existing interaction techniques for the digital library can be categorized in **direct manipulation**, **selective manipulation**, **zoomable user interface** (ZUI), **dynamic exploration** and **interactive exploration**. The rest of this chapter presents the related work in each of these major categories.

2.2 DIRECT MANIPULATION

Direct manipulation is considered a significant HCI design improvement over the command-line systems prevalent prior to the GUI and WIMP era [22, 39, 96, 84, 100, 144, 171, 179, 188, 190, 194]. The goal at the time was to eliminate the painful process of memorizing cryptic type-in commands. By conforming to a consistent design, it allows the average user to learn a system quickly. Direct manipulation is perhaps the most implemented interaction technique for the digital library, including Perseus. Instead of typing in commands at a prompt, users can interact with screen GUI objects (lists, links, buttons, menus, drop-down, etc.) directly using the **object-action interface** (OAI) [183] paradigm. A typical page design displays a region with text and images. There may be a region reserved for showing tree hierarchy of files and folders. A click (or an equivalent action) on a node in the tree hierarchy is equivalent to a 'go to' command. Direct manipulation for hypermedia environments allow users to navigate from one article to another by clicking on links among documents, bookmarks, citations and annotations. User interactions involve command-like WIMP actions such as scrolling, clicking and branching.

Research efforts in direct manipulation concentrate in improving techniques in specifications, screen design, querying, retrieval and zooming [4, 15, 22, 198]. A good design in manipulation generally encourages the consistent use of set standards or protocols (placement, color, etc.). The major advantage of the direct manipulation is its short learning curve and established user base. The major disadvantage of the direct manipulation is its possible lack of potential for future development. Since its fundamental design originated from the desire to replace a command-driven system, user interactions involve command representation rather than information representation. This is perhaps why it is very cumbersome to find summary or semantic information in existing interfaces for the digital library. Because direct manipulation may pose difficulties in representing an information space, it is not uncommon to temporarily 'break away' from a direct manipulation system and shift to a different interaction style, such as a panorama or a virtual tour.

2.3 SELECTIVE MANIPULATION

The interface design of selective manipulation is similar to direct manipulation, but it goes beyond interaction with screen objects. The screen layout also employs the desktop metaphor with GUI and WIMP interactions. However, in selective manipulation, user selections may not equate to a direct command to perform a task. Instead, user selections shift the visual representation and move the users within different data layers. For example, a family tree shows the relationship (an arrangement) of a collection of people (information context). Users may select a node (a focused member of a family) in the tree to access another data layer (personal information sheet of the selected family member). The tree shape can also change (rearrangement) in reaction to user request (select a new ancestral root). Therefore, the interactions of a selective manipulation concentrate in the arrangement and rearrangement of the information space. To that purpose, the screen display generally suggests selection and encourages the users to make the right selections. Popular visualization styles for selection manipulations include variations of the trees [40, 62, 63, 122, 168], maps [18, 59, 121], books [35, 207], charts, graphs or hybrids [134, 105, 180]. The screen layouts generally contain a breadth-oriented (as in maps) or depth-oriented (as in trees) hierarchical visualization. The book metaphor can allow both breadth (lots of pages) and depth (links on pages) selections.

2.3.1 INTERACTION WITH TREES

The 2D tree graph is one of the most popular geometric visualizations for node-link, networking and connectivity. Tree graph can show the entire tree relationship (context) while maintaining easy access to node details (focus). Tree structures can be top-down, right-left, or center-out (star, radial). Interaction techniques of the trees generally involve arrangements and rearrangement of the nodes. Navigation in a hierarchical structure involves moving from one node to another along the existing hierarchical links in the structure. The tree is very effective for a small-sized hierarchy. When the size of a hierarchy becomes large, navigation becomes challenging because embedded in the complexity of tree is the same focus versus context problem that many information visualization systems have attempted to address.

In 1996, Lamping and Rao proposed the **hyperbolic browser**, a radial tree, as a focus+context technique for visualizing large hierarchies. [26, 122, 123, 124] It was originally developed by Xerox PARC and much work has been done on it by its spin-off company, InXight. A variation of the hyperbolic visualization is licensed to the Microsoft® **SiteMap** [56]. The web site of the Library of Congress also employs the hyperbolic browser with over 20 top-level sub-trees branching out from the root node. [45] Hybrids of the hyperbolic browser have been a major thrust of research in the display of large data sets. The focus of the hyperbolic browser is generally the center of the tree while the context is represented by the complete tree. Relationship is given by the edges of the nodes. The neighboring nodes are directly linked in certain contextual relationship to the focal node. The hyperbolic browser provides a focus+context view of the data by arranging a center distortion and placement of the entire tree structure all at once. If users drag a node to the center, the tree structure rearranges to reflect the hierarchy of all remaining nodes with respect to the center node. Many different variations of the hyperbolic browser have been implemented to reflect branching, linking, literature and document relationships.

Another example of the tree structure is the **Neigbourhood Explorer** which employs a star tree for house hunting in London. Nodes are images of houses that meet a certain searching criteria. [7] Radial scales are used to represent attributes such as price, area, rooms, etc. The problem with this schema is repetition because a house can be represented in more than one scale path (i.e. once on the 'price' path, once on the '# of bedrooms' path, and so on). The advantage is that users can drag a house to the center and the star arrangement is automatically adjusted to reflect the deviation from the attribute scales. Another work along the same grain but with a stronger zooming capability is the dynamic **HomeFinders** [206] by Williamson and Shneiderman.

The tree structure, most particularly, the hyperbolic browser, is a good tool for reflecting branching relationships. Therefore, if the connectivity and branching factors are well balanced, the tree can be an effective tool for representing information such as table of contents, citations and library collections. Beyond that, the tree structure has severe limitations in representing a large document space because most library literature does not translate well into a hierarchical relationship.

2.3.2 INTERACTION WITH MAPS

The map metaphor is an effective visualization for representing distributions, as well as topological, regional and classification relationships. In 1989, Kohonen first proposed the concept of a **self-organized map** (SOM) [113, 114] in which a complex information space is represented by colored regions with borders. The regions are constructed by unsupervised learning algorithms. Each region contains a group of documents similar to each other in some ways. Maps such as the SOM are abstract in nature. Similar methods have been adopted in the **ET-Map** created by a research group at the University of Arizona in 1998 where neural network algorithms are used to generate a map of an information space based on more than 100,000 documents about entertainment on the Web. [45] Other pioneering work in this area was done by Lin in 1997. [131, 132]

Maps can be a degenerative case of the 3D geographic paradigm, which includes information landscapes, terrains, topologies, panoramas, virtual tours, murals, walkthroughs and storytellers. [42, 83, 102, 103, 110] Some of the classics in maps are the contemporary maps of **London underground transportation system** (‘**The Tube**’) [72]; Florence Nightingale’s diagram of the death rate in 1858; Harry Beck’s maps of Cholera deaths in London’s Soho district in 1845; Charles Joseph Minard’s map [189] of the “Loss of the

Napoleon's Army" in 1812 and its contemporary improvement, the **SAGE** [173], by Carnegie Mellon University (CMU). Various other versions of the map metaphor permeate digital libraries in a wide range of forms, such as **literature map**, **domain map**, **document map**, **contour map**, **semantic map**, **article map**, **citation map**, **data map**, **science map**, **semantic map**, **feature map**, **context map**, and the like [18, 59, 114, 186]. The names generally suggest the functions and these maps can be a hybrid of the abstract maps or tree-like hierarchies as in the **tree map**. [105, 108, 180]

2.3.3 INTERACTION WITH BOOKS

The book metaphor is perhaps the most intuitive paradigm for the digital library. Implementation of the book metaphor includes virtual books, electronic books, bookshelves, reading rooms, and other supporting metaphors of the reading space. [35, 84, 207] The title, thickness, volumes or even wear-and-tear of a book collection can convey contextual information about the subjects. The **Workspace of Web page** by Card et al. in 1996 at Xerox Corporation allows the users to interact with 3D pages, books, shelves and desktop space in a data landscape. The **WebBook** and the **Web Forager** present a faithful rendition of the book metaphor with interactions in opening and closing of a book, turning and flipping the pages. [35] The **workspace for document** by Risch et al. in 1997 showed pages distributed in a 3D space with wire-grid geometrics for abstract data, such as a scatter graph. **Lifestream** by Freeman and Fertig in 1995 showed event pages layered on top of each other;

the X and Y axes are used to see a topmost page in detail and the Z axis is used to show a timeline of events.

When used effectively, the book metaphor is strong in suggesting actions in reading and learning. Unfortunately, its deficiency can sometimes outweigh its advantages. In real life, bookbindings afford such actions such as selecting the book, opening, closing and flipping the pages. In cyberspace, users do not need to pick up the book, carry it to a desk, and open it in order to read it. Such interactions can be meaningless and inefficient. It is also superfluous to show page flipping because a click can change the text. The animation of page-flipping action does not improve reading. Such actions do not contribute to information absorption and actually increase system overhead. The graphic elaboration of a book may compete with text space. Due to its limitations, the book metaphor is frequently downgraded to an image backdrop on which two blocks of text are displayed. User actions can be downgraded to direct manipulation with buttons for actions such as 'next page,' 'previous page,' 'home,', 'end' or even scrolling. Nonetheless, the book metaphor will continue to be popular due to its long establishment in human civilization.

2.4 ZOOMABLE USER INTERFACE (ZUI)

The term zooming user interface (ZUI) is relatively new, although similar ideas and techniques have been used in computer systems for remote sensing image processing and geographical information systems for years. [128, 159] A ZUI system must allow users to bring focus to the area of interest freely and easily, across various levels of detail. [156] Zooming involves clicking or dragging on a region to bring focus to it. In **fractal** (planar) **views**, focus is brought about by techniques in enlargement, magnification, clarity, highlight, colors and so on.

In **attention-warped displays**, various techniques in **distorted views** have been developed to zoom to the focal region. [37, 38] Zooming displays can be **time multiplexing** or **space multiplexing**. In time multiplexing, the same screen space is shared and users zoom to the focus or the context at different times. In space multiplexing, two or more screen displays are available at the same time. In **spatial zooming**, the context and the focus represent the same information, as in the **document lens** or the **magic lens**. [23, 170] In **semantic zooming**, the context is represented by symbols and the focus is the actual information. An example of semantic zooming is the maps. In **zoom-and-replace**, if the users click on an area of interest, a different display replaces the original display as the zoomed-in view.

If zooming is a used as a technique in amplification or clarification of the same data, especially in space multiplexing, then it is most often called **overview+detail**. In overview+detail visualization, the nature of information can be the same; whereas in focus+context visualization, the nature of information in the detail document and the overall context should be different and additional. Techniques in focus+context involve time multiplexing of semantic zooming of contextual information. A well-implemented ZUI can achieve an impressive visual effect. However, distorted ZUI can expensive to implement. If an image is too strongly distorted, it can severely damage legibility. Therefore, ZUI should be used discreetly for in actions that require reading.

2.4.1 ZUI IN FRACTAL VIEW

One representative ZUI example is **Pad++** [19, 20, 151]. It is a multiple scale display tool jointly developed by a number of universities, including NYU and UMD. A ZUI designed with Pad++ would enable the users to zoom across a wide range of granularities. [19, 20, 21, 151] For example, a ZUI design of the cosmos can allow virtual visitors to explore our universe in various levels of detail. Users can zoom in and out infinitely: from an overview of our solar system, the earth, London to a model of a human body, the human's heart, blood cells, and so on. The **flip zooming** shows the context of the documents in thumbnails of different sizes. [24, 93] Users zoom to a document in full page, which assumes the characteristics of the book metaphor such as page flipping. The magic lens is similar to the fisheye but with no distortion of the text. In a way, the magic lens can work like a movable information filter [191] or a magnifying glass [204]. The fisheye view can be applied to SOM where the area of interests bulges out. The zoom enlargement technique is frequently used in image browsers. [2, 118, 120, 184, 206] The prominent ones in this category are the collection of 'finder' projects, such as the **PhotoMesa**, **PhotoFinder**, **FilmFinder**, and **LifeLines** [155], where drag-and-drop is used as a strategy for labeling photos, films, videos, etc. The context is given in a large colored starfield graph and the focus is given by a zoomed enlargement of the active image. PhotoMesa is called a **zoomable image browser** by the HICL (HCI Lab) at UMD. [91, 104] The zoom areas are different directories where the photo images are stored. The users navigate the space of images. The design tries to stay clear of the direction manipulation paradigm and there is no management of scrollbars, menus, pop-up windows. PhotoMesa has its own image layout algorithms called **Quantum Treemaps** and **BubbleMaps**. Work at UMD also tries to synchronize multiple views or **snap-together** visualization on a 2D screen. [147]

2.4.2 ZUI IN DISTORTED VIEW

Distorted focus+context visualizations are also known as attention-warped displays. Much work has been done in the arrangement and rearrangement of special visions that depart from the conventional desktop metaphor. The **bifocal display**, originally developed by Spence and Apperley in the early 80's, places the entire context in a rolling strip with a colored area of focus. [189] The concept presented in the bifocal display concept was invented long before a working model could be built. In 1986, Furnas presented the **fisheye view** that became a popular paradigm for the display of large hierarchies. [12, 68, 69, 175, 176] To overcome the focus versus context problem, the display is distorted in spherical vision, presumably the way fish see. In a fisheye view, the center of the orb in **degree of interest** (DOI) is the focus defined by user action in clicking or dragging. Local details of the focal point are enlarged while the rest either remains the same or becomes proportionally squeezed. In 1991, Mackinlay et al. developed the **perspective wall** [136] using the same concept in 3D perspective where users see well-proportioned frontal views but distorted side views. Similarly, the **movie wall** of Indiana University is a wall for displaying frames in a movie in much the same way as the perspective wall. The document lens is similar to the fisheye but uses planar perspective with two vanishing points.

2.5 VIRTUAL 3D

Virtual 3D is less explored as a solution domain to the focus versus context problem. Virtual reality can be the ideal environment for spatial tasks, virtual museums, historic buildings and archaeological sites. [11, 80, 83, 89, 146, 152, 166, 174] Convincing 3D designs can put a heavy demand on computer resources. [13, 78] Animation is also frequently employed as part of the user-interface design for a virtual environment. [9, 42, 43, 58, 169] A true 3D system is considered to be immersive. However, since immersiveVR generally requires special equipment beyond the reach of most users, desktopVR or fish tank VR is becoming more common in the digital library. In desktopVR, the 3D world coordinate system is mapped onto a 2D screen coordinate system; therefore, desktopVR is sometimes referred to as **2½D** or **3D-like**. The screen space is often a symbolic representation of the 3D space. Simple applications may involve a virtual tour or panoramic views. More complex models may involve walkthroughs, storytelling and educational learning. [3] Interactions in desktopVR are inventive uses of the existing WIMP technology in symbolic immersion. For example, ↑ key moves forward, ↓ key moves backward, ← key turns left, → key turns right and so on. Whether they are 2D½, 3D-like or 3D, there are roughly two types of virtual applications: dynamic exploration and interactive exploration. [8, 106, 140, 126, 203] The virtual scenes in a dynamic exploration generally do not change in context. Users can learn about the environment by navigating dynamically in the fixed scenes. The virtual scenes in an interactive exploration are more life-like to reduce learning curve by using carry-over knowledge from real world experiences.

2.5.1 DYNAMIC EXPLORATION

Dynamic exploration seeks to create a data environment where the users can freely explore the information space. Using a spatial metaphor, an application in dynamic exploration encourages the users to move toward the area of interest. User location defines the focus of the contextual space. Nearly all 2D paradigms such as trees, maps and books can be upgraded to benefit from the added real estate in a virtual environment. For example, the **3D cone tree** is a 3D counterpart of a 2D tree graph. [168] On a 2D screen, there are limitations to how large a tree can be squeezed into the display area without losing legibility. In order to overcome such limitations, a variety of distorted or warped spaces have been proposed and developed. The 3D cone tree takes advantage of the spatial freedom in 3D. Many more nodes can be displayed with reduced clutter. However, more complex user interactions are required to navigate the tree structure. The 3D cone tree can also be used in conjunction with other visualization tools in 3D. For example, the prototype of **Cat-a-Cone** in the Yahoo forager was built on the basis of 3D cone-tree and animation from the Xerox PARC **Information Visualizer**. Another important component used in Cat-a-Cone is based on the WebBook. The WebBook and the Web Forager are designed to provide a more intuitive way for users to access interconnected information. Category labels are only displayed if they match the current document displayed in the focal book in the foreground. This information workspace also includes a bookcase. The user can search the collection documents and save the results of a search as a book. If a book is not currently in use, it can be shelved in the bookcase.

2.5.2 INTERACTIVE EXPLORATION

Interactive exploration goes beyond dynamic exploration and requires more duplexing between the virtual world and the users. Interactive exploration generally adopts an advanced interpretation of a variation of spatial metaphor. A convincing virtual environment demands realistic modeling of a believable world. Animations, tour guides and avatars can be added to the scene to further increase the sense of interaction. Unfortunately, an interactive exploration system can put a heavy demand on computer resources and programming overhead. Therefore, it is not uncommon to see virtual prototypes using simple geometric objects (i.e. cubes, spheres, etc.) as symbolic stand-ins of the more complex counterparts. For example, a virtual robot can be constructed using a sphere for a head, a cylinder for a neck and a stack of rectangular blocks for the body and arms. The reason for the simplicity is generally driven by the economics in reducing the cost of implementation.

Many interactive exploration applications present a room-like setting. The design concepts allow users to behave in a virtual world analogous to a real world. The **VRoom** [202] is a project at the MITRE Corporation that uses the spatial metaphor to create a multi-walled (i.e. 6 or 8), customizable room where multiple users can go in and obtain information. Categories of information are distinguished on different walls in a very realistic room setting. A ‘sand table’ in the center of the room can convert 2D maps to 3D topologies. An intelligent and informative avatar can assist users with additional information. Microsoft© has a similar application called the **TaskGallery** that handles internal project progress. A research project developed in Germany, in cooperation with the University of Rome ‘La Sapienza,’ called **VIRGILIO** [200] uses a visual metaphor for video and music database visualization. It allows users to interact with structural information in virtual reality. Users can navigate across different level of data in different scenes, from buildings to rooms, and access drawers that correspond to a musical CD and related information.

The virtual 3D world can also be abstract or geometric as in the **Selective Direct Manipulation (SDM)** [46, 49]. SDM consists of layers of XY planes. Each plane contains a table of 3D color-coded data columns. The tools include SDM handles and widgets that users can interact with to modify the location of the plane (shift operations). However, abstract 3D worlds must be carefully executed. Otherwise, the system may fail to provide real-life carry-over behaviors and the users may not know how to maneuver in a complex geometric scene. Consequently, an interactive exploration system may degenerate to a dynamic exploration system if the users simply navigate in the virtual environment completely unaware that they can interact with the objects in the world in a special way.

2.6 LITERATURE SUMMARY

The focus versus context problem has been explored from many different angles. A major research area concentrates on the development of more effective tabular, graphical or hierarchical display of library contents. There has been much research in the development of better query, browsing and access techniques for 1D (e.g. list), 2D (e.g. table XY), 3D (e.g. space XYZ) and N-dimensional (e.g. space and time XYZT) data contents. [47, 71, 111, 117, 131, 150] Another major research area concentrates in more effective design and layout on a very small 2D screen. [27, 41, 84, 94, 115, 116, 177] With affordable graphics, there is an increased research interest in visualization and interaction in virtual 3D spatial management and design. [10, 32, 36, 60, 61, 66, 86, 132, 138, 186] Temporal or abstract information visualization generally employs techniques in knowledge transfers afforded by an effective metaphor. The most successful metaphor in use today is the desktop metaphor. Many other metaphors have been invented to help convey complex information. Some popular metaphors for the digital library are books, geographical, maps, and distorted visions. In general, it can be safe to say, that geometric representations can be more effective in bring context but not focus and that metaphoric representations can be more effective in bringing focus but not context. In geometric representations, the context tends to be the primary display that loses the screen space to the secondary display, i.e. a selected focus. For example, a tree displays primarily the hierarchical relationship. When a node is selected, the secondary display (e.g. a page of information about the node) can block the visualization of the tree in part or in whole. In metaphoric representations, the focus tends to be the primary display (e.g. an object in a room, a page in a book, etc.) relative to the contextual background (e.g. room, book, etc.). It is difficult to maintain the overall context (e.g. all rooms, table of contents) while navigating from one focus to another. In short, the focus versus context problem remains to be a challenging research topic.

As the research forges on, the digital library continues to benefit from technology improvements in many fields in computer science, including human-computer interaction, information visualization and virtual reality. [148, 160, 149, 161, 172] As the demand for better user interface increases, the overhead in design and implementation increases proportionally. More research effort has been dedicated to improvement in direct manipulation, selective manipulation and ZUI than in dynamic or interactive exploration. Virtual 3D systems continue to be experimental and have not gained a steady user base. The number of applications of virtual reality is endless – but they all hinge on whether or not such systems can be built. While the technology is still immature, we must have reasonable expectations.

CHAPTER 3

PROPOSED SOLUTION

In this chapter, we propose a set of new interaction techniques for the digital library as a solution to the focus versus context problem. Our solution methodology abandons the conventional research approaches in better screen design, placement, retrieval and access of web pages. Instead, our new approach seeks to examine the true nature of the information items and bring focus to them by artistic 3D reconstruction in a large virtual scene. We seek to organize the data foci in virtual 3D arrangements designed specifically to bring out the context of information. We have designed a context-rich environment with data-rich focus using 3D focus+context interactive visualization. We have implemented an information environment where users can interactively visualize focus without losing context. We believe that a 3D information space allows users to learn better and faster because they are immersed in the environment. The sense of immersion relates to real-life learning experiences as users 'walk about', 'look around' and 'pick up' visual information from the world around them. We have chosen virtual reality as our implementation platform for its advantages in added virtual real estate, 3D realism and the sense of presence. We seek to validate our approach by conducting usability testing to evaluate the benefits of our proposed solution. The design concepts proposed in this chapter demand high-quality implementation in desktopVR and immersiveVR. Our solution concepts do not seek to replace the existing paradigms in digital library. Instead, we believe that they can be alternatives or additions to the current approaches.

3.1 DESIGN CONCEPTS

Our 3D world has been meticulously preserved in literature. In essence, 3D context is converted into 1D text describing our world and the events that happened. The writings are supplemented with illustrations and photographs. For example, an excerpt of lengthy text, with a plot map of the intersection, is used to describe the London Fire of 1765 (Section 4.4). We can say that information has been downgraded from 3D to 1D or 2D in a form of **context crippling**. During the early days of computing machines, printed materials are converted to text documents. Pictorial information has been digitized at a slower pace due to the slower development of computer graphics. With the advent of the world-wide-web, vast quantities of physical documents in all subjects have been and are being converted everyday into electronic forms in full graphics. Ironically, it is the downgraded versions that are being digitized in the digital libraries. Even though 3D graphics are becoming commercially available, the images are still 2D as part of the legacy to the photographic era. Again, in a certain perspective, we can also say that the 3D data are being downgraded in a form of **data**

crippling. For a long time, restoring the 3D nature of information may have been prohibitive due to the overhead in virtual reality systems. More importantly, though, is the fact that many of us have become so accustomed to the **information crippling** that we have a tendency to design and populate a 3D scene with 2D objects. For example, many systems show 2D document sheets of 1D text floating in 3D space. Great pain is taken to show 2D images as billboards that face the users at all times. Part of the reason was to reduce implementation overhead, but the driving force is a lack of recognition in restoring the 3D nature of information. Once again, we can call this effect a form of **virtual crippling**.

In this thesis, we recognize these **crippling effects** as a major deterrent to visualizing and interacting with information in the digital library. Our design concepts seek to remedy the crippling defects by redefining 3D focus and context. We recognize that the importance of developing effective visualization techniques lies in our far superior ability in image recognition, as compared to our relatively poor ability in text parsing. Therefore, we strive to restore the 3D nature of information context in virtual reality. We realize that a large context landscape can be slow in real-time rendering. However, current computer speed is not our primary concern. Instead, we hope to show a new frontier in lightweight interactions using visualization techniques in virtual reality. We seek to bring out the information buried in digital library documents and display them in virtual 3D. Merely showing the documents in 3D is not good enough; we need to be able to visualize and interact with the information represented by the documents. We believe the crippling effects are inherent in the problem of focus versus context. Our proposed solution seeks to immerse the users in a computer-generated environment that is reflective of a real world where the data populate the scene in a meaningful way. Virtual reality can be a powerful platform for visualizing and interacting with information because the environment simulates the way users already learn about the world around them in real life. The transfer of knowledge can allow users to absorb spatial information with little or no training in virtual reality.

Our proposed solution consists of three parts. First, we seek to redesign 3D focus and context in a way that uses interactive visualization and lightweight interactions in virtual reality. Second, we seek to demonstrate the feasibility and functionality of our designs by building high-fidelity prototypes. We have built three major prototypes to fulfill this part of our mission: the London walkthrough (Chapter 4), the Greek coin catalog (Chapter 5), and the Greek vase museum (Chapter 6). We argue that it is almost intuitive that users who can 'visualize' the 3D content (i.e. London, coins, vases), as compared to 'reading' about it in 1D text, should learn about the same focus+context information in less time and with greater accuracy. Third, we seek to demonstrate the advantages of our proposed solution in a usability evaluation of the Greek vase museum. We will present the results and analysis of our experiment in Chapter 7.

3.2 INTERACTIVE 3D FOCUS

The main focus of digital library visualization is the documents. Obvious upgrade of digital library documents to virtual reality would be an implementation of the 2D book-and-page metaphor. The problems with this approach are cluttering, occlusion, and illegibility. The viewport is still confined to the front-most 2D plane while the rest of the document space is hidden behind the pages closest to the reader. There is only the appearance of more reading power and we are not exploiting the extra real estate available in the virtual domains. The cluttering can be controlled by categorization, spatial organization and use of colors, but the 3D space can become too busy. The occlusion problem can be mitigated by the use of transparency, but at a cost of reducing legibility even more. Legibility can be enhanced with magnification by proximity, such as fisheye and table lens, but only in a limited way.

3.2.1 3D TRUE NATURE PRINCIPLE

We propose a virtual environment with a sense of computer realism that adopts a **true nature principle** where we create '*a 3D context landscape populated with 3D data foci that are modeled as faithfully as possible to their counterparts in the real world*, *given available source information*.' It is insufficient to build a realistic-looking 3D world. Instead, the objects in the virtual environment must be representative of some real counterparts in the real world. Examples of data focus adhering to the true nature principle are accurate renditions of a locale (city, streets, room, etc.) or of data artifacts (coins, vases, gems, sculptures, etc.). The goal of applying a 3D true nature principle is to provide a data-rich environment where users can benefit from studying the objects in the virtual world. For example, if a designer creates a virtual room with imaginary furniture, the users cannot learn from the environment.

However, if the designer creates a room with ancient Greek furniture true to nature, then the users can learn by immersion, as compared to reading text passages or viewing drawings of the furniture in digital library. If the data is representative of a physical object in the real world, we can base the 3D modeling on the real counterpart. However, if the data is abstract, adhering to the true natural principle requires information research.

Simulated virtual scenes can even offer some advantage over real-life experiences. A visitor to London, for example, can only absorb the sights and sounds of his immediate surroundings, with no sense of the city as a whole. Similarly, a student can tour a museum but cannot touch or rearrange the artifacts. Such restrictions need not exist in a virtual world. If the data is representative of entities of a different scale in time and space, users may not be able to visit or interact with the real counterparts at all. This is the case where users may not visit the galaxies in context, even though the constellations are replicated in visualization models and we know they exist. Such is also the case where interactions of visualization at a cellular level must be translated in scale and demand high-level skill sets.

Building 3D models adhering to the true-nature principle can be time-consuming and painstaking. However, the benefit is worth the effort. The more accurate they are, the more likely that they will enjoy longevity in reuse. 3D models adhering to the true-nature principle can become portable components. Interface designers with access to these models can populate a more meaningful virtual environment in less time. This allows them to focus more on their design concepts than on graphical implementation. Limited resources in computer power and man-hour can be obstacles in building virtual environments that adhere to the true nature principle. User-interface designers should tap into techniques in rapid component prototyping (Appendix A). Given the sophistication of computer technology today, crude mockup of sophisticated worlds (using basic primitives) should not be the only alternative if the reason for such crudeness is the lack of resources.

In our three prototypes, we have constructed the London buildings, Greek coins and Greek vases true to the nature of the selected information available in Perseus. We have extracted 242 buildings from the original maps (Appendix B.1) and saved them as JPEG (JPG) files. We then used different versions of the JPEG files in various renditions of the London walkthrough. FIGURE 3.1 shows the process of extracting a sample JPEG file, i.e. a building numbered t_lb085 (Section 4.3.2).

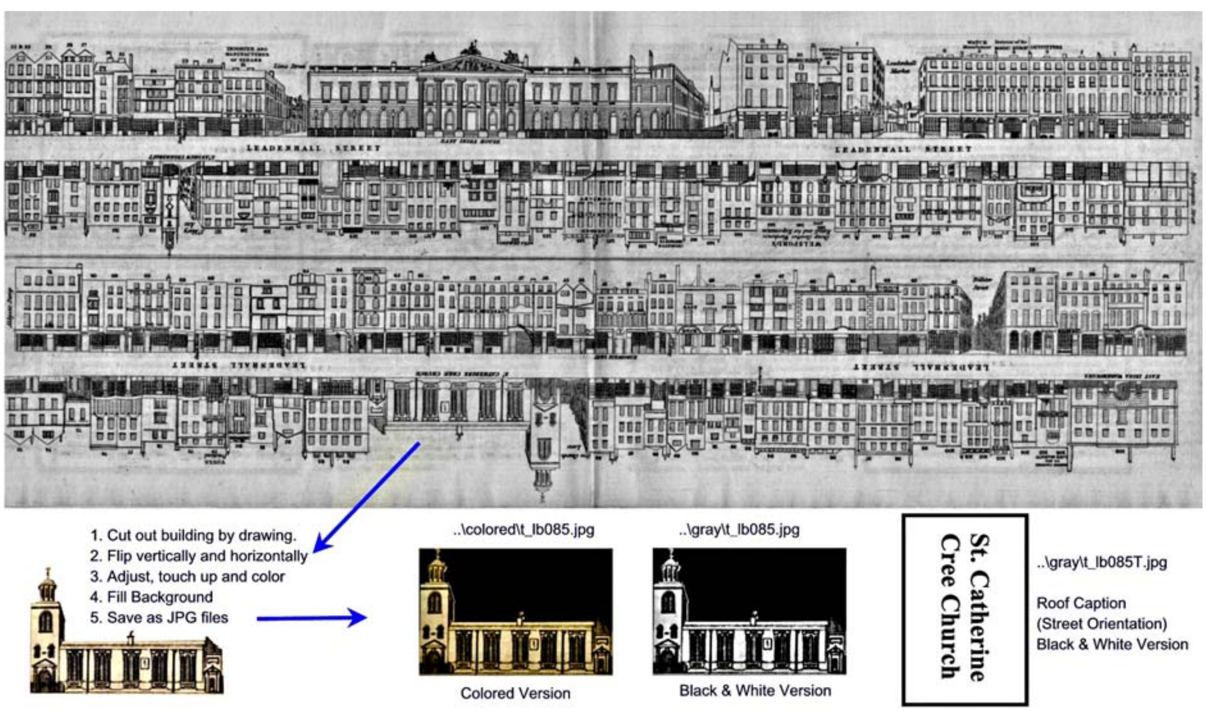

FIGURE 3.1 EXTRACTING JPEG TEXTURE MAP OF BUILDING t_lb085

FIGURE 3.2 shows the location of building t_lb085 in a rendition of the London walkthrough.

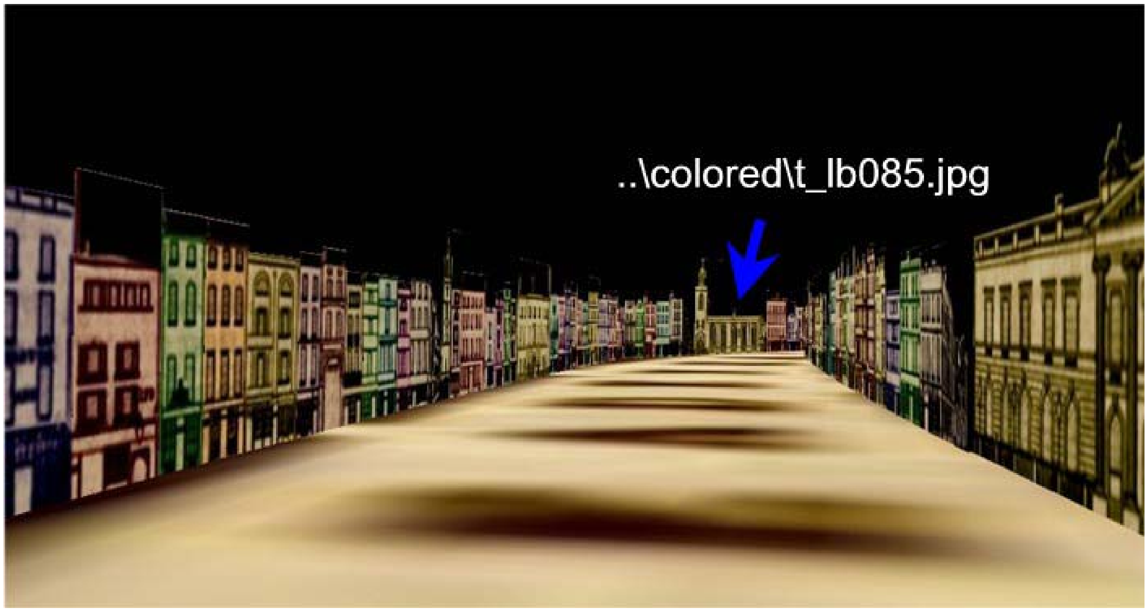

FIGURE 3.2 BUILDING t_lb085 IN THE LONDON WALKTHROUGH

The JPEG files used for the coins and vases are modeled analogously. FIGURE 3.3 shows some samples of 3D Greek coins and 3D Greek vases modeled true to nature, by texture-mapping JPEG files onto 3D geometric models of the data artifacts.

FIGURE 3.3 SAMPLE 3D GREEK COINS AND VASES

To populate the virtual scenes in our three prototypes, we have constructed 531 3D models (242 buildings, 132 Greek coins and 157 Greek vases) true to nature. The 3D models are independently stored and can be used as component objects in other projects.

3.2.2 3D INFORMATION ART (INFOART)

Next, we propose InfoArt, which is '*the artistic rendition of 3D data focus, using educated decisions, as a way to visualize missing contextual information*.' The difference between the true nature principle and InfoArt is the artistic license with informed decision. InfoArt goes beyond the true nature principle and calls for the talents of graphic artists who can draw and paint in 3D with the absence of accurate metrics (a skill set that may take a while to develop). A data object that merely applies the true nature principle is mechanical. Designers follow a set of known metrics in modeling them. However, when information is incomplete or purely abstract, InfoArt seeks to fill the gap by providing an artistic but informed rendition. The graphic artist is able to do so either by extensive research or by relying on expert inputs. For example, it is possible to reconstruct a 3D sculpture of the head from skull remains. Some information calls for InfoArt. For example, different virtual renditions of the mythological tales based on researched data can allow users a different learning experience. We propose a resurrection of the 3D nature of data artifacts using informed judgments. Virtual worlds can greatly benefit from artistic interpretations. Why should we be looking at 2D pictures of real objects when 3D models can be cheaply built given today's computer technology? Once built, the 3D data become a permanent part of a digital library collection. The longevity of information art and the contextual data they represent is thus worth the effort in making them. In reality, advanced InfoArt techniques must be implemented with expert knowledge and professional accuracy.

In our prototyping, we used a crude form of the InfoArt technique in the modeling process of our vase museum. FIGURE 3.4 shows a London vase with a catalog number of 1920.12-21.1. The right side shows a semi-transparent HTML page triggered by the proximity sensor as the users approach the vase in the virtual scene. The vase has no available picture. However, the caption on the left tells us that the vase is a 'Pyxis' (a canister). We can then generate a grey model of a Pyxis in the virtual scene for this vase.

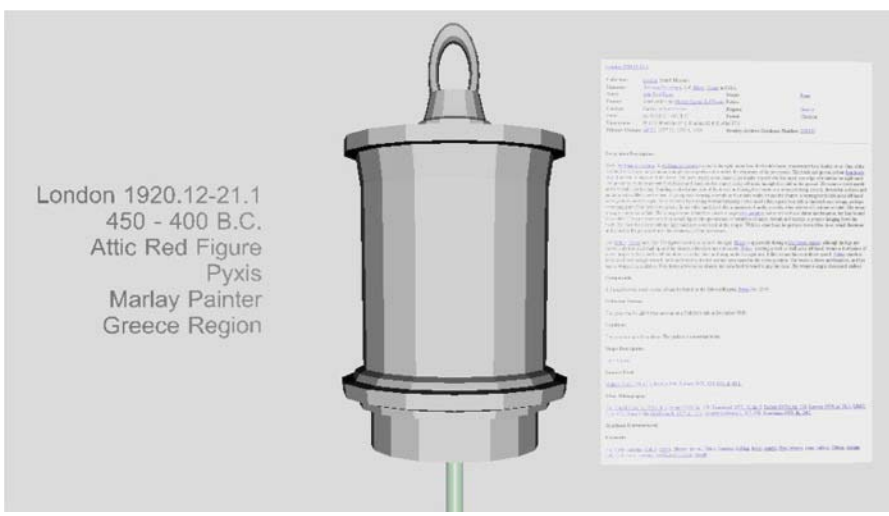

FIGURE 3.4 LONDON VASE 1920.12-21.1

Some vases have only partial pictures. In this case, we can determine the shape of the vases from the text description. We can also determine its looks from the neighboring vases with similar catalog numbers. For example, London vase B 280, shown in FIGURE 3.5, has only a partial picture on the HTML page. However, we can determine its shape and texture map by its immediate neighbor, London B 226, shown in FIGURE 3.6. Both vases are neck amphora and attic black figures made in the Etruria region.

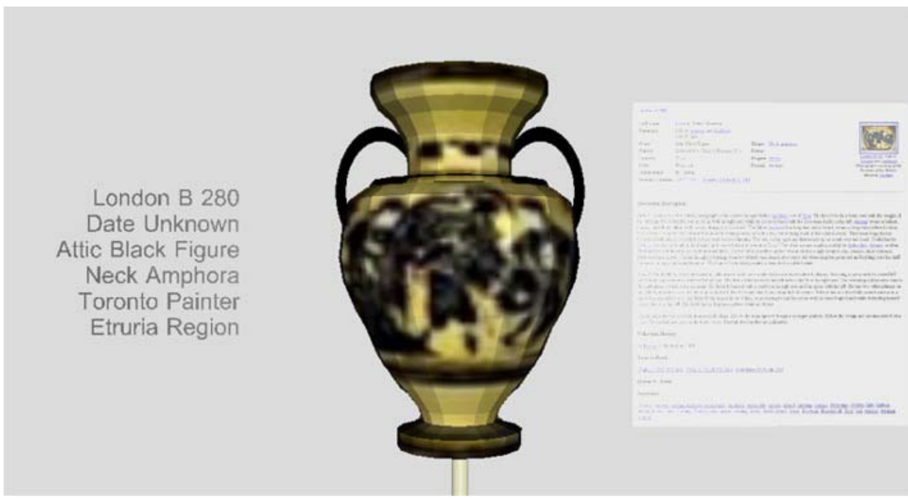

FIGURE 3.5 LONDON VASE B 280

FIGURE 3.6 LONDON VASE B 226

In a crude fashion, when certain pieces of information are missing, we try to construct the 3D models based on informed decisions. Without the creativity in InfoArt, the virtual scene would be incomplete and fragmented. In our vase museum, we inform the users about our use of the InfoArt technique in a subtle way. For example, they know intuitively that 3D grey models in the scene have no pictures. Furthermore, as they approach the HTML pages of the vases, they will see a partial picture on certain vases.

3.3 INTERACTIVE 3D CONTEXT

Intuitively, it is harder to model intellectual information than quantitative data. Before a 3D model can be built, the information structure must be determined and the most appropriate paradigms must be selected. Sometimes the information structure is self-evident. For example, to build a walkthrough, it is best to reconstruct the geographical locale in virtual 3D as close to the real scene as possible. After all, one of the primary goals is to eliminate the need to travel. In this case, the paradigm is easily determined. However, sometimes the information structure can be elusive. For example, representing 4D (XYZT or space and time) is expensive in virtual reality. A virtual environment is defined by the XYZ axis and is an ideal platform for represented 3D data. However, humans live in 4D space. It is more difficult to represent London in the 1840's as compared to what it looks like today. The two time settings can be modeled separately in different virtual environments; advanced interactive techniques in scene morphing could achieve the effect of time changing. It is even more difficult to change part of a virtual scene without losing the sense of immersion. For example, it is very difficult to show the construction history or resident turnover on a particular street. Refreshing part of a virtual scene can take some time and the virtual environment may appear to be ‘frozen.’

Limitation in graphics and computer speed also restrict visualization and interaction with a large data set. Unfortunately, technical feasibility rendered many complex interaction techniques to be wonderful in theory and unusable in practice. In order for information to be presented in context, the semantics of the documents need to be extracted from a digital library. The underlying context of the documents extracted must be interpreted and an appropriate visualization determines the document modeling required, which is materialized in a virtual 3D scene. The modeled contextual landscape allows users to explore and navigate in 3D space using natural (i.e. walking, reading, etc.), natural-plus (i.e. flying, morphing, etc.) and lightweight (i.e. proximity, level of details, etc.) interactions.

3.3.1 3D INTERACTIVE WALKTHROUGH

We propose using 3D interactive walkthroughs as the virtual space to represent location information. The major difference between existing walkthroughs [34, 83] and the interactive walkthroughs we propose is that users can visualize a data-rich world and focus on a building while maintaining the context of our information space. The 3D interactive walkthroughs we propose go beyond existing techniques in browsing, image viewing, quicktimeVR, panorama or simple virtual tours that are implementations of image-swapping techniques. In an interactive walkthrough, the world must be created true to nature, allowing users to navigate and explore embedded layers of information. As users approach a region, more information about the location may appear in the scene. For example, an interactive walkthrough is an ideal platform for storytelling, time morphing and historical renditions. In time morphing, the scenes can gradually change, in whole or in part, to reflect changes in the information space. For example, a London street can be time-morphed from 1845 to 2003. Unfortunately, due to slowness in rendering speed, our prototype excludes morphing.

3.3.2 3D INTERACTIVE MUSEUM

We propose using a 3D interactive museum as the virtual space to represent data artifacts. In an interactive museum, objects on display are placed in a room-like setting where users can explore on their own. Using lightweight interactions, user interests are determined by their proximity to an object. An interactive museum differs from the traditional virtual museum in many ways. The data artifacts displayed in an interactive museum must follow the true nature principle in restoring the object in 3D information art. Displaying 2D images of the data artifacts is not sufficient because that downgrades it to a 3D image gallery. The layout of an interactive museum should resemble an actual museum layout (e.g. a specific exhibition in the Boston Museum of Fine Arts). The layout of the museum can be imaginary, abstract or geometric, but it must be created in an effort to convey appropriate information.

The most important advantage of an interactive museum over a virtual museum is the visitor's ability to navigate and interact with the data artifacts on display. For example, users can pick up an object, examine it and put it back. As an area of interest is approached, more information can be displayed in the scene. Users may request more information about the displays. There can be a curator in the scene to help users with their needs. The exhibits can be arranged and rearranged on demand to highlight various aspects of the information space. For example, users may ask the 'museum' to sort the collections based on certain criterion – something they cannot do in a real museum. The complexity of an interactive museum is defined by its realism and the level of interactions.

3.4 INTERACTIVE 3D FOCUS & CONTEXT

Excessive branching and connectivity can distract users from their original intent. Current digital library implementations implicitly require users to know beforehand exactly what they are looking for. If they fail to provide the crucial key words, they will be bogged down in a hierarchy of futile searches. The human ability to absorb information from the environment is lost on a 2D screen and blocked by keyboard entry. Users are unaware that while they are searching for 'information' (the right goal) pertaining to specific topics, the computer is querying the documents by matching values in the fields or text and returning a 'list of names' (wrong results) for the documents. The discrepancy between search purpose and system results is intrinsic in design and cannot be corrected by better searches or screen designs. The information is at best grouped by categories and keywords. Our solution proposes a new way of organizing 3D layouts. Although we organize our 3D worlds realistically, we also blend in the advantages in geometric representations. The two focus+context designs we are presenting are the **3D map overlay** and the **3D interactive graph**, both contain an overtone in geometric representation over realistic renditions. In the map overlay, buildings (focus) are modeled to maximize information and placed to reflect the orientation of the streets, plots and contours. In the 3D interactive graph, data points (focus) are placed in a geometric layout in such a way as to reveal quantities of information.

3.4.1 3D MAP OVERLAY

We propose the 3D map overlay as the layout of 3D walkthroughs (Section 4.1 and 4.4.2). Many of the current 3D renditions of virtual environments are created in an ad-hoc fashion. Either the models are created without any consideration of placement accuracy or a photographic image is used as a stand in. Buildings and street orientation should not be placed randomly or expediently. Instead, they should be placed in proper coordinates in agreement with the scale of other maps that define the locale, such as street maps or plot maps. Designers and implementers should have some realistic anchors to structure the walkthrough orientations. In a 3D interactive walkthrough, the location information must be accurate by certain metrics. Generally, geographic information can be defined by a set of maps. Enforcing a map overlay technique puts a heavy demand on the designer and implementers; however, the resulting context is more meaningful to users who are navigating the virtual environments.

Chapter 4 presents a prototype implementation of the overlay technique in an interactive London walkthrough. The 3D models of buildings are reconstructed from a set of building map drawings. The width and depth of the buildings are reconstructed to match the sizes on a set of plot maps. The buildings then form long streets that are bent to conform to the contours from a different set of street maps. In essence, we are combining three types of maps in one scene: building maps, plot maps and street maps. Users can walk around and see the true bending of the streets. From an aerial view, they can see the plot maps conforming to the city streets. A major advantage of the map overlay technique is that the placement of the data is in context with the metaphor. The arrangement and rearrangement are natural and intuitive. Smooth navigation is important to the sense of realism in an interactive walkthrough. 3D zooming may be required to allow users to quickly move into an area of focus. If users are interested in a particular building or street, they can click (in desktopVR) or 'touch' (in immersiveVR) it and the system will fly in and land on eye level. 3D zooming can be simulated in camera views.

3.4.2 3D INTERACTIVE GRAPHS

We propose 3D graph visualization as a variation of the map overlay. In a map overlay, the ground is a geographical or topology map which defines the placement of 3D physical objects in the scene. The map overlay must be observed from an aerial view of a geographical setting. It may be difficult to see a map overlay in an eye-level view without a location indication, which may downplay the effect of lightweight learning. Navigation must allow users to fly or change camera view quickly. In an interactive 3D graph, the information landscape is the geometric chart and users are immersed in the data that are populated data artifacts (Section 5.2, Section 6.2). For example, the Greek coins can be data points (since they are round) in a 3D graph by reign (i.e. X-axis), collection (i.e. Y-axis) and period (i.e. Z-axis). Any wall structure (China's Great Walls, the Berlin Wall) or even river flow can form a line graph (since the information nature is linear) through XY topology and Z height/depth. Buildings can be columns in a table chart. The streets in the 3D interactive London walkthrough can also be converted to a line graph. We can place different streets in their natural contour and compare the buildings, the cost of construction, and so on.

The basic guideline that satisfies a 3D graph is that the data points must be true to the nature of the information represented and the users must be able to interact with the data in the graph in a certain way for the purpose of learning more about them. In our prototypes, we presented several sample implementations of the 3D graph concept. In the interactive Greek coin catalog, we presented two collections of coins in a 3D scatter plot. As the users navigate the plot, they see 3D models of the ancient coins 'floating' in midair in 3D coordinates by year and collection. A bird's eye view of the 3D scatter plot shows a degenerate 2D scatter plot. In the interactive Greek vase museum, we display a collection of ancient vases on transparent columns with very light tints of colors. Each color represents a different region that the vases are made. We place the vases in the room by year and ware (type of glazing). As the users navigate the vase museum, they see the vases clustered by year and wares. A bird's eye view of the room shows a 2D scatter plot by year and ware. A far eye-level view shows a histogram of the column colored by regions. In short, as the users navigate in our three 3D prototypes, they can see the information we have extracted from Perseus in different virtual graph views.

3.5 PROTOTYPE OVERVIEW

The interests of this thesis lie in the interpretation and representation of a new set of interaction techniques for the digital library. Our solution must be able to handle large datasets of textual, non-textual and abstract data, such as images of artwork, maps, artifacts and other such information. The design concepts of our 3D worlds call for realistic reconstruction of large datasets with believable 3D objects. Building the high fidelity prototypes described in this thesis requires significant resources in computer graphics, memory management, device and virtual reality programming. We also propose a portable prototyping environment where immersive equipment (i.e. head-mounted display, trackers and gloves) can be outfitted on a laptop. We have built three significant prototypes to demonstrate and validate our design concepts. The prototype environment is described in Appendix A. Our first prototype is an interactive London walkthrough created from a selection of map overlays. Users know about the London fire and want to learn more about how the fire affected the intersection they are visiting. The prototype changes to reflect the London fire in the 1840's. As users visit the intersection, they become interested in a museum they saw and decide to go in. Once they are in the museum, they become attracted to some Greek exhibitions, specifically the Greek coin and vase collections. To this effect, the second prototype is a Greek coin catalog and the third prototype is a Greek vase museum. The coins and vases are replicated in virtual 3D models. Instead of displaying the coins and vases in rows of shelves and glass cases, the coins are displayed in a 3D scatter plot populated by data coins and the vases are displayed in 3D data clustering degenerative to a 2D scatter plot. No direct command is ever given in our 3D worlds; instead, the virtual environment can glean the users focus by physical proximity and virtual navigation. The next three chapters discuss the prototypes in greater details.

CHAPTER 4

3D INTERACTIVE LONDON WALKTHROUGH

The 3D interactive London walkthrough (London walkthrough) is a virtual street intersection modeled from three 2D drawings of London maps. In this chapter, we will present the design rationale and implementation issues in the creation and renditions of the London walkthrough. We will show the current approach of learning about the London maps in Perseus. The current approach fails to provide a sense of context or mental model. The new approach of learning about London involves exploring a 3D model of the intersection. Users will have the added benefits of lightweight interactions while they navigate the virtual streets. As an example of practical application, we will use Perseus and the London walkthrough to learn about a London fire that affected the intersection in 1765. We will first present a text passage found in Perseus that described the fire (Section B.2). The text passage makes little or no reference to the original maps of the intersection. We will then present several screen captures from our London walkthrough reflecting the fire. We believe that the interactive walkthroughs afford a spatial learning experience not available in existing 2D screen space.

4.1 INTRODUCTION

Users interested in London street scenes can view various versions of the Tallis maps in Perseus. A sample of the Tallis maps is shown in FIGURE 4.1.

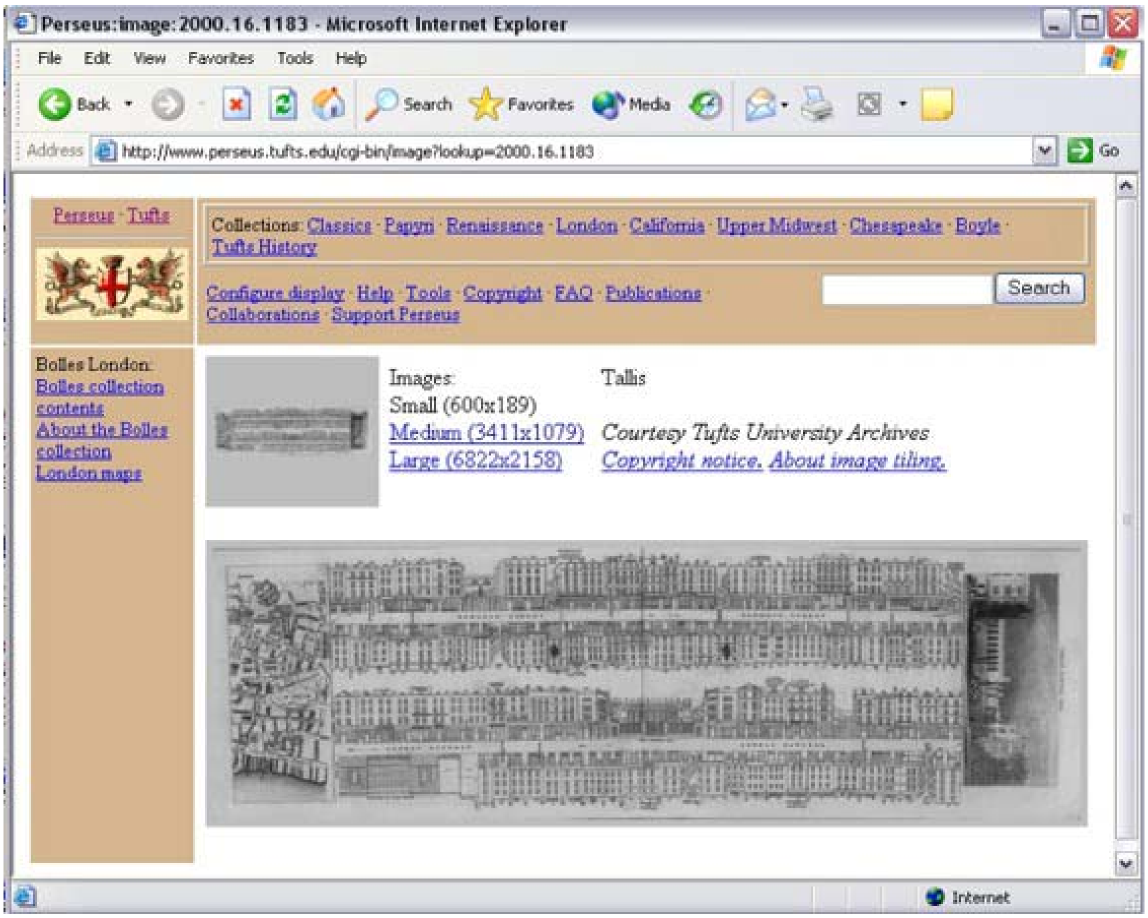

FIGURE 4.1 SAMPLE TALLIS MAP IN PERSEUS

These Tallis maps are the creation of **John Tallis** who was one of the most popular cartographers of the 19th century. He was renowned for his accurate and attractive maps of all world areas during the Victorian Age. Principal engraver and mapmaker to John Tallis was **John Rapkin**, renowned for his extraordinary artistic and engraving skills. In 1851, John Tallis & Co. issued the **Illustrated Atlas**. These are among the most popular nineteenth century maps available to the collectors. Although most of their maps contained a limited amount of color, collectors of his maps were known to have commissioned colorists to tint the beautiful vignettes. The maps are very detailed, with copious topographical and political information presented with excellent clarity. Perseus contains copies of many of the famous Tallis maps in its archives, the Bolles Collection. [30]

The organization of the Tallis maps poses a serious problem in focus and context. Unlike other images or text, there is no apparent context to the maps stored in Perseus. There is also no focus because of the poor legibility of the online maps. Users need to choose between focus (a detailed section of a map, such as a building) versus context (the set of relevant maps). For the Tallis maps, the challenge starts from defining focus and context first before anything else can be done. If users are mission driven, locating relevant maps is not as difficult as understanding and digesting the maps found. As these maps are scattered without context it is up to the users to figure out the relationship among them. It is not practical to fit all relevant maps on one screen. If one map is shown at a time, the users will lose the contextual information of the other maps. Regardless of the orientation and display technique, it is very difficult for most users to form a mental model of the information represented in the maps. Therefore, users are compelled to learn about the maps in an ad hoc manner. They tend to gloss over the maps in their browsing experience without appreciating the geographical connection among the maps. Even if they were informed of such correlation, most users would not invest the time and energy to analyze the maps or attempt to form a mental model.

4.2 DESIGN RATIONALE

VR can offer a superior 3D information space for visualizing the maps. The benefit we seek to add is to allow the learning of map information by 3D focus+context visualization of the virtual environments where users can walk about and explore on their own. A building is the smallest focus and the map overlay is the context. Streets and intersection are larger foci. Organization of the buildings is by streets and is inherent in the real-life city metaphor. As users dynamically navigate in the 3D world, they can concentrate on one building without losing the context of the entire intersection. In an interactive walkthrough, the context is inherent in the street scene where users can see all the buildings of the intersection in one glance. Occlusion is natural to the metaphor and can be informative. For example, it is natural to have a building block other buildings and the streets disappear into the horizon. Users can move around and get different views of the intersection. Users can focus on a building, walk closer to it and examine it in greater detail, without losing perspective with the rest of the buildings in the intersection.

To demonstrate our design concept, the London walkthrough we have built is a virtual reality prototype of a 4-way intersection reconstructed from three Tallis maps found in Perseus. [53] The four streets modeled are: **Bishopsgate Street**, **Cornhill**, **Gracechurch**

Street and **L**<u>**L**</u>**eadenhall Street**. The intersection is named **BCGL** by the first letter of each the four streets in alphabetical order. Drawings of the buildings in BCGL are spread over three disjointed 1840's maps from which 3D texture maps are extracted to model the virtual environment. There are 242 building models in the intersection. The set of original maps used to model the BCGL intersection can be found in Appendix B.1. FIGURE 4.2 shows a screen capture of a center corner. Four more screen captures are shown on the next two pages. FIGURE 4.3 shows a scene of Bishopsgate Street, looking from the end of the street toward the center. FIGURE 4.4 shows a scene of Cornhill, looking from the end of the street toward the center. FIGURE 4.5 shows a scene of Gracechurch Street, looking from the end of the street toward the center. FIGURE 4.6 shows a scene of the Leadenhall Street, looking from the middle of the street toward the center.

FIGURE 4.2 CENTER SCENE IN LONDON WALKTHROUGH

FIGURE 4.3 BISHOPSGATE STREET

FIGURE 4.4 CORNHILL

FIGURE 4.5 GRACECHURCH STREET

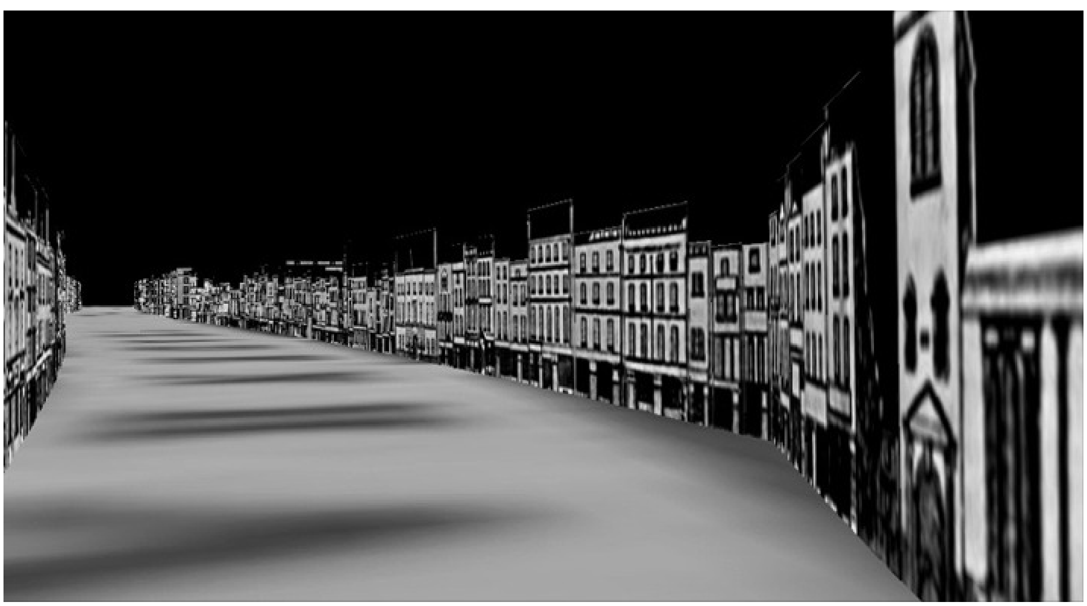

FIGURE 4.6 LEADENHALL STREET

FIGURE 4.7 shows the map overlay visualization that merged the 3D spatial quality of three types of 2D maps: street maps, plot and building maps. The buildings are aligned in exact proportion to match the contour on the street maps.

FIGURE 4.7 LONDON WALKTHROUGH MAP OVERLAY

As the users get closer, the building plot sizes become labeled rooftops. Instead of seeing a 90-degree quadrant-like intersection, users can see the actual curvature of the streets. The 3D walkthrough allows users to maintain focus and context at all times. The geographic paradigm 'takes the users there' in the hope that the users will form a cognitive map or a mental model with its basis in landmark knowledge and development in survey knowledge. Currently, users have to access different maps to obtain different types of information. In an interactive walkthrough, one or more 2D maps can be combined using the 3D map overlay technique. In the London walkthrough, the rooftops of all buildings reflect the plot map. In a bird's eye view, the users see the plot maps of the intersection superimposed on the street map of the city. Given the original set of three maps, we believe that users may not be able to identify the intersection from three seemingly unrelated maps. Even if they were able to come to that conclusion or be told of that fact, it will still take them much more time to examine the maps and come up with the mental model of how the intersection should look in 3D than they can by navigating the London walkthrough. The goal is to allow users to learn better and faster about the intersection in 1840's London in focus+context interactive visualization. Instead of browsing images of the maps, users can learn about the intersection by lightweight navigation.

4.3 IMPLEMENTATION ISSUES

London is a very old city with intertwined streets, alleys and back roads. A 90-degree intersection is in essence nonexistent. With the help of an archivist, the BCGL intersection was chosen because it best conforms to a 4-way intersection. Data completeness is also a major factor in a true nature reconstruction. Although missing buildings can be represented as 'holes' or grey areas, a fragmented scene is not as desirable. Simple boxes with 6-facet texture maps are used to model the buildings, as compared to the more expensive architectural meshes. Buildings form streets, streets form the BCGL intersection and final touches, such as bending, sky, grounds, lights and cameras, complete the London walkthrough. Finally, the entire scene is exported to a VRML file to be rendered by the browser. The following subsections discuss some of the implementations issues encountered in the making of BCGL. The implementation is empirical. The size of the intersection pushes the limits of typical desktop workstations in loading speed and graphic ability. In the London walkthrough, there are 283 objects with 1,888 vertices and 2,832 faces. The physical memory usage is between 271.4 MB (minimum at frame 0) and 511.5 MB (maximum). The virtual memory usage is between 218.6 MB (minimum at frame 0) to 1,506.3 MB (maximum). It takes approximately 00:00:25 seconds to produce the view shown in FIGURE 4.7 at the maximum frame rate of 100 (see Appendix A, Performance). The users are fixed in a walking mode, with a speed of 2.0 and a height of 75 units per second.

4.3.1 SIZES AND SCALES

Determining the sizes of the buildings in the virtual world can be more of a challenge than it first appears. Intuitively, the sizes should match that of the real world. Not only are the real-world coordinates historically correct, but the dimensions would seem to offer higher sense of realism if the viewer is to walk down the streets in a virtual environment Unfortunately, the initial intuition turned out to be optimistically naïve. A small 100x300 JPEG (JPG) file stretched to the size of its corresponding real building will become pixellated. The world is so huge that viewer sees a blotchy mess of inky dots. The effect is more like blowing up a JPEG file 800 times in a paint program than anything realistic. The scenes can also appear to be pixellated if users are too close to the building. Since John Tallis had given great care to the proportions of the buildings in his drawings, we decided to preserve his scales in the reconstruction of the buildings. The problem with using the scales on the Tallis maps is the discrepancies in the scanned-in proportions of related maps. For example, BCGL is reconstructed from three maps, but the scales of two adjoining streets, Leadenhall and Bishopsgate, are different. When the initial version of the buildings is reconstructed in virtual reality, buildings reflected different heights. The solution of this problem in later versions of BCGL reconstruction is to match the offending scales to the majority of the scales used in the scene. Another problem encountered is matching street maps to plot maps. However, not all plot maps of the buildings in BCGL can be found, and depth dimensions are missing from these maps. For the sake of expediency, the building depth is set to be the width. Bending of the BCGL streets also distorts depth information. When the straight 90-degree, 4-way intersection is bent to fit the contour of the aerial street map, the intersection conforms to the layout shown on the original maps.

4.3.2 LEVEL OF DETAILS

Ideally, nothing short of a faithful rendition of the architectural details is acceptable. Unfortunately, such reconstruction is impossible because the Tallis maps provide only the fronts to the buildings in simple 2D drawings. It is impossible to construct an accurate historical walkthrough without more collaboration from other disciplines. The baseline of the buildings is horizontally aligned, but the roof line is not. To faithfully render the scene, the buildings ought to be modeled in meshes showing the contours of the different roof heights, dormers, chimneys, etc. But this is an expensive proposition and must be abandoned if the walkthrough prototype is to be run in real time. Instead, simple box primitives must be used to save space and speed up process. This is because if a box exports to a VRML97 file of 400 bytes, its mesh equivalent can be about 7,400 bytes. A cheap solution to this problem is to fill the skyline and negative spaces of the JPEG texture map files with black to outline the different heights and extrusions of the rooflines. When viewed in 3D with a default black background, the rooflines of the aligned boxes appeared to be contoured without the use of complex meshes. This is a sacrifice of the architectural detail for navigation speed.

In addition, all buildings are modeled individually as basic box (rectangular) primitives. Buildings modeled individually allow a large intersection to be warped to fit the curvature reflected on the original street maps in bird's eye view. This is because a minor bending applied to selected buildings offers an overall appearance of a subtle curving of the streets. Had the building been modeled in the largest possible sections of 512x512 images, the bending of the streets could appear sharply angled. The JPEG files required to texture map the buildings are created by cutting out the building images from the Tallis maps (Appendix B.1) using a drawing program that can support cutting by drawing. Since the buildings' contour is hardly ever a perfect rectangle (most are hand-drawn with lots of imperfections), a background color matching the building is applied to eliminate the white spaces that may show through in between buildings when the streets are reconstructed. Crooked buildings are pulled straight by a small rotation in the XY plane. Upside-down buildings are flipped vertically and horizontally. The building images are touched up as necessary to enhance contrast and to add tints. All the JPEG files used in a scene are stored in a separate directory. The same scene can be directed to load texture maps from different JPEG directories. There are three renditions of the BCGL JPEG texture map files stored in three separate directories: 1) ../gray – contains the black and white images of the building; 2) ../colored – contains a tinted version of the buildings with arbitrary colors; and 3) ../fire – shows the fire-damaged buildings. The BCGL intersection contains 242 building images that must be individually cut out and touched up. Each building has a roof caption JPEG file that shows the building number and name. Roof captions double the number of JPEG files used in the intersection. Other supporting images required by the scene are prepared analogously. It takes about 500 JPEG files, or ~50 MB of images, to model the 242 buildings in the BCGL intersection. A building image is named by the street it is on, the bounding street, and it sequential number in the original Tallis map. If a number is missing on the original map, it is interpolated from the two nearest buildings. For example, t_lb085.jpg (FIGURE 3.1 in Section 3.2.1) stands for the texture map JPEG file of the 85th building on Leadenhall, bounding by Bishopsgate, on the Tallis map that it is extracted from, and t_lb085T.jpg is the image file for the roof caption.

Material libraries are indispensable in creating a complex scene where there are hundreds of JPEG files to be used as texture maps. To facilitate later export of the scene, the material libraries and all its associated JPEG texture maps files are stored in one directory. Fancy materials such as stuccos or marbles should be avoided because they appear to be specks and blurs when viewed through a head-mounted display (HMD). Contrasting colors should be used delineate bordering objects because the HMD does not support as many colors as a standard monitor. A light-brown object next to a mustard-colored object may appear to be a blob of yellow. In general, scenes viewed in a low-budget HMD tend to be lighter and shinier. To offset the brightness, scenes in the London walkthrough are created to be darker and gloomier.

4.3.3 LOADING ISSUES

The most painful hurdle in building our London walkthrough is the sheer size of the resulting file, which poses a major challenge for many loaders. A smaller scene with just a few buildings will run much faster, but such simplicity is not realistic for most real street scenes. As mentioned in earlier subsections, the architectural details and roof line contour have been sacrificed to speed up navigation. But each building has a unique frontal image and that cannot be compromised. Many browsers recommend light use of texture maps in real time, and the BCGL intersection has far too many buildings and texture-maps. Ideally, texture maps should be small and applied sparingly because they can increase the file size, slow loading, and cause a potential memory burn. Unfortunately, every building in the Tallis maps is different. Each must be modeled faithfully to match the original maps.

An alternative solution is to model a large box texture mapped with a huge box with numerous buildings. Unfortunately, some browsers will not display the texture maps if all the vertices of an object are outside of the current viewport. Therefore, a large flat surface must contain enough vertices in them so that at least a few vertices can be seen at any possible viewer position in a virtual environment. In essence, this is equivalent to breaking the large box into smaller boxes or creating meshes. Another necessary violation of loading efficiency is the used of multi/sub-object materials, which texture-maps the building image to the proper facet of the building box. Unfortunately, they are exported as separate objects in VRML97 because VRML97 does not support more than one texture map per object in library export. This problem can be overcome by more efficient programming in future work.

Another issue in morphing of a large intersection is speed. It is possible to simulate large-scale morphing by physically switching to a different scene. However, such abrupt behavior discounts the effect and defeats the purpose of morphing. In practice, loading a large scene can take a long time and the sense of continuation and immersion can be compromised. Another problem in morphing is that in order to time-change one object to another, the underlying meshes must be structurally the same with congruent geometric vertices in order to morph from one copy to another. The drastic reallocation of land and changes in building plots can cause severe implementation problems because the underlying meshes cannot be the same. Therefore, dynamic scene morphing cannot be implemented and rendered in real time, cost-effectively, in our prototypes.

4.4 APPLICATION : THE LONDON FIRE OF 1765

The London fire of November 7th, 1765 started in the early morning and rampaged for approximately six hours (3 a.m.–9 a.m.) before it was brought under control. A complete text passage describing the fire can be found in Appendix B.2. Although the streetscape must have changed from 1840s to 1765, for demonstrational purpose, it is still possible to trace the same buildings that were burned. A rendition of the London walkthrough is used to reflect the London fire on that date. In this section, we will compare the current approach in using Perseus and the new approach in the London walkthrough.

4.4.1 CURRENT APPROACH

The users can learn about the London fire by searching and browsing related sites in Perseus. FIGURE 4.8 shows the Perseus site that describes the fire. The complete text can be found in Appendix B.2.

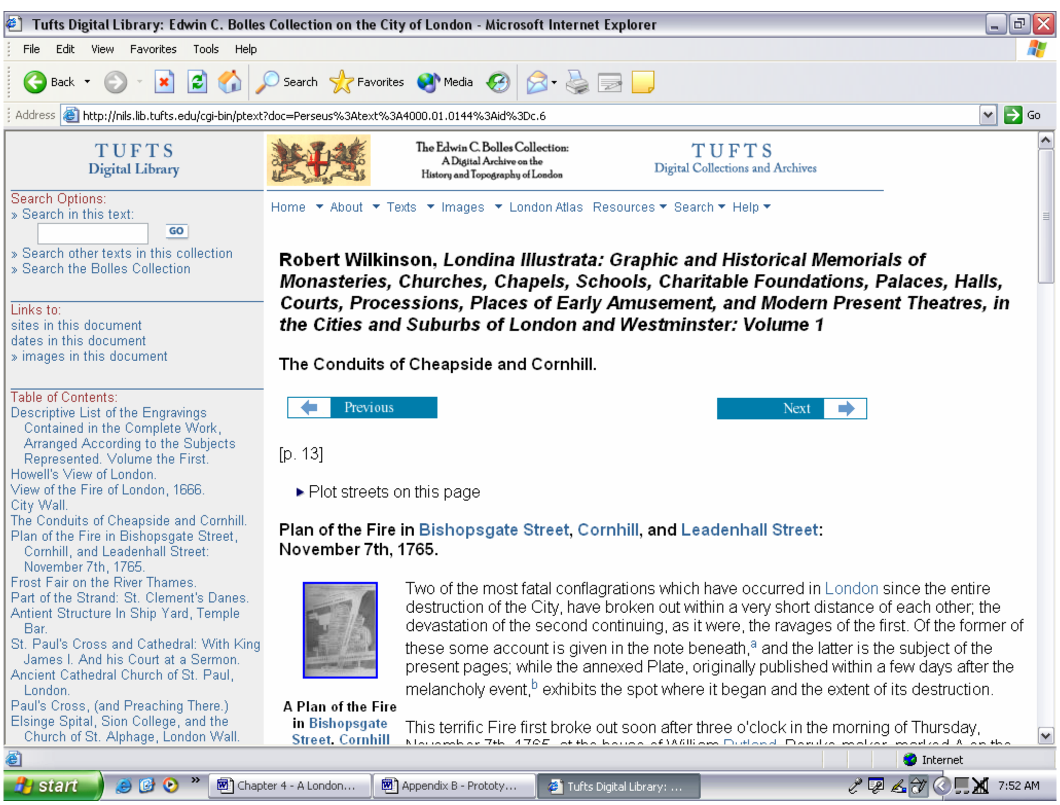

FIGURE 4.8 HTML PAGE OF LONDON FIRE IN PERSEUS

4.4.2 NEW APPROACH

FIGURE 4.9 shows a screen capture of the London walkthrough. The fire damage is shown by the use of colors. The undamaged buildings are shown in its original grey. The damaged buildings are shown in various shades of red. The amount of damage is shown in different intensities of red. The first three buildings that triggered the fire are numbered in cyan. The fire path is indicated by the direction of the arrows on the buildings. In black and white mode, the virtual environment is seen as different shades and intensities of grey.

FIGURE 4.9 LONDON WALKTHROUGH OF FIRE DAMAGE

This rendition of the London fire is a replica of the same BCGL prototype described earlier with the same geometries but different set of texture maps loaded for the buildings.

There are three renditions of the BCGL intersection: the black-and-white version, the color version, and the fire version. In the color version, random colors are chosen to highlight the buildings. FIGURE 4.10 shows a screen capture of a bird's eye view of the fire. It is can be shown that fire damage is mostly confined to one side of the intersection (upwards in the map). The south-west wind steered the fire away from the intersection. The part of Bishopsgate to Threadneedle Street was completely burned. Three angles of the BCGL were burned. The corner of Cornhill and Gracechurch was spared.

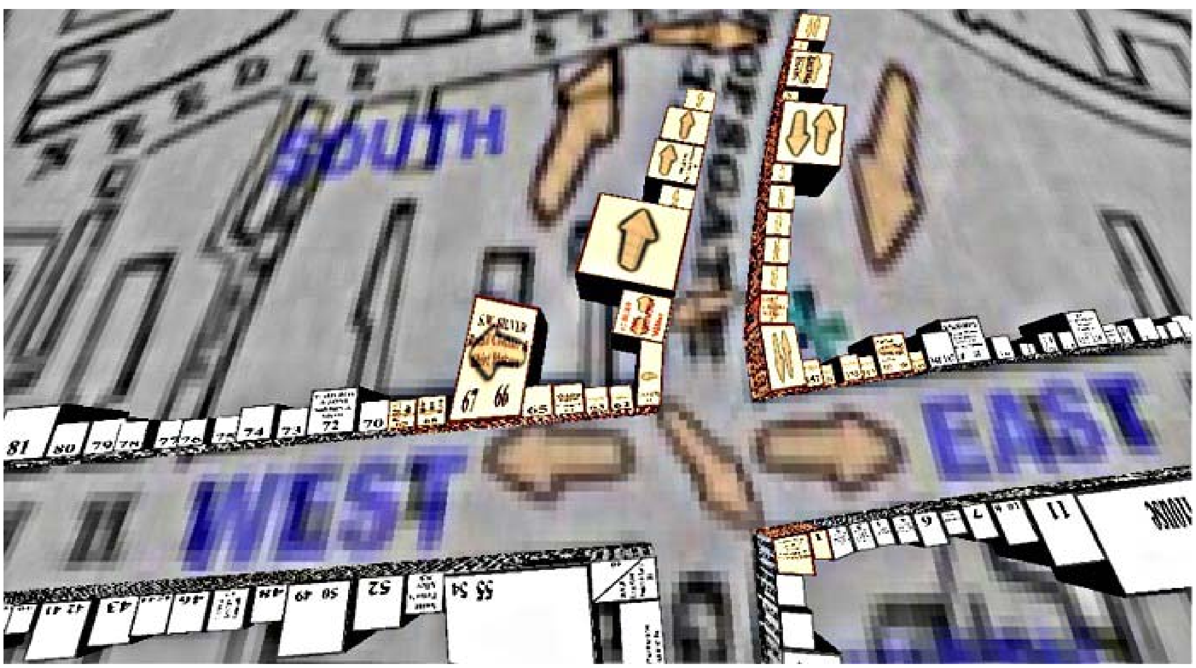

FIGURE 4.10 MAP OVERLAY OF FIRE DAMAGE

It is possible to combine the current approach with the new approach. Storytelling can be incorporated in the London walkthrough to show users the events of the fire. The text pages describing the plan of fire shown in FIGURE 4.8 can be converted to floating text page.

The pages can be transparent or semi-transparent, with scrollbar widgets to control the flow of text. Many of the web book or web page techniques can be applied to storytelling in 3D. [35] Eye tracking can be used to trigger additional passages. If the users stare at a building for a certain amount of time, more text about the building can appear in the scene. As they look away, the text can disappear. The users can read the text as they walk around in the scene. A 3D text menu can allow users to 'touch' a keyword on the page and the virtual scene can change accordingly and 'take' them right to the corresponding building in the scene. FIGURE 4.11 shows a very crude implementation of the 3D text menu.

FIGURE 4.11 3D TEXT MENU OF FIRE DAMAGE

4.5 SUMMARY & FUTURE WORK

The new visualization implemented for the London walkthrough is the map overlay, where several 2D maps are manifested in 3D space by mere placement and orientation. A new interaction technique implemented is a 3D text menu. Users are allowed to dynamically explore the environment and interact with the information represented in the world in a lightweight fashion. The buildings are reconstructed true to the nature of the original map drawings. The visualization is 3D geographic and the interaction style is 3D interactive exploration, where users are like tourists visiting the virtual locale. Given today's computer technology, some of design concepts we have proposed may not be achievable in real time. However, fragments of high-fidelity mock-ups can be implemented to offer a glimpse of what the 3D information space could look like.

The concept of a 3D interactive walkthrough can add immediate value to current image browsing applications. There are many reasons to why 3D interactive walkthroughs are less prevalent. First, maps data are all over the place, disjointed and incomplete. Second, such data are difficult to obtain, requiring the collaboration of a cartographer, historian or an archivist. Third, most maps are 2D line drawings and do not transform well into impressive 3D views. Without an impressive appearance, most graphic artists do not choose historical map renditions as the top choice of their modeling projects. Fourth, it takes a lot of effort to plow through the maps and build the models. Fifth, affordable computer speed is too slow to handle a large intersection, much less a complete city.

Future work in 3D walkthrough include improvements in 3D modeling, inclusion of architectural information, multi-layered 3D spaces, direct programming to reduce exported file size, providing database storage for building structures, enhanced zooming and selection in desktopVR. More work can be done in the area of immersion, tangible user interface (TUI), 3D reading, location indicator and navigational control. Interactions in an interactive walkthrough are navigational, proximity-driven and lightweight in nature. The information landscape is data-rich. If users get closer to a building, it can be that they are interested in learning more about the building. More information (such as residents and costs) can emerge in the scene. When they walk away, the extra information can disappear. We can have world within worlds [65] in a multi-layered scene, users can walk into a building and see the interior of the building. The interior of a building can be a true historical rendition, or it can be an information room with data about the building, or a combination of both. Avatars can be added to the scene as either tour guides or hosts of the buildings. Home and business owners can serve as hosts or tour guides. Users can receive 4D data, or XYZT (XYZ and time), reflecting time change. Users can request to see the same intersection 100 years before or 100 years after. Users can see how the buildings change through time of the day. This is not currently implemented in the BCGL walkthrough, but it can be simulated easily by switching scenes. Users can issue a vocal command, touch a building or select from a 3D text menu. ZUI and 3D navigation can allow users to zoom in or peter-pan in bird's-eye view.

CHAPTER 5

3D INTERACTIVE GREEK COIN CATALOG

The 3D interactive Greek coin catalog (coin catalog) is a 3D scatter plot populated by selected coins from two collections of Greek coins in Perseus. In this chapter, we will present the design rationale and implementation issues of the coin catalog. We will compare the current approach versus the new approach of learning about the Greek coins. The current approach involves browsing the corresponding web sites in Perseus. The users can see pictures of the coins and read descriptions about them on the HTML pages. The current approach promotes the users to choose between the focus (one coin) versus the context (the collection). It also promotes the users to access one collection at a time or one list with the coins mixed in alphabetically, regardless of collections. The new approach in learning the coins is by exploring a 3D coin catalog we have built. The users can learn about the Greek coins by immersing and exploring a virtual room where the coins are the data points of a 3D scatter plot. Both focus of the coins and context of the collection are maintained at all times. The users can easily distinguish the similarities and differences of the various collections. Virtual 3D can also afford direct interactions, such as grabbing, with the 3D coins.

5.1 INTRODUCTION

Perseus contains a wealth of coin collections from various sources. The coins are identified by the collection and their catalog numbers, as shown in FIGURE 5.1 (i.e. Dewing 2423 represents the Dewing collection, catalog number 2423). The users can access the coin catalog in Perseus from its home page and other major sites. The users can select a coin from a list in the Perseus Coin Catalog by clicking on the corresponding catalog number.

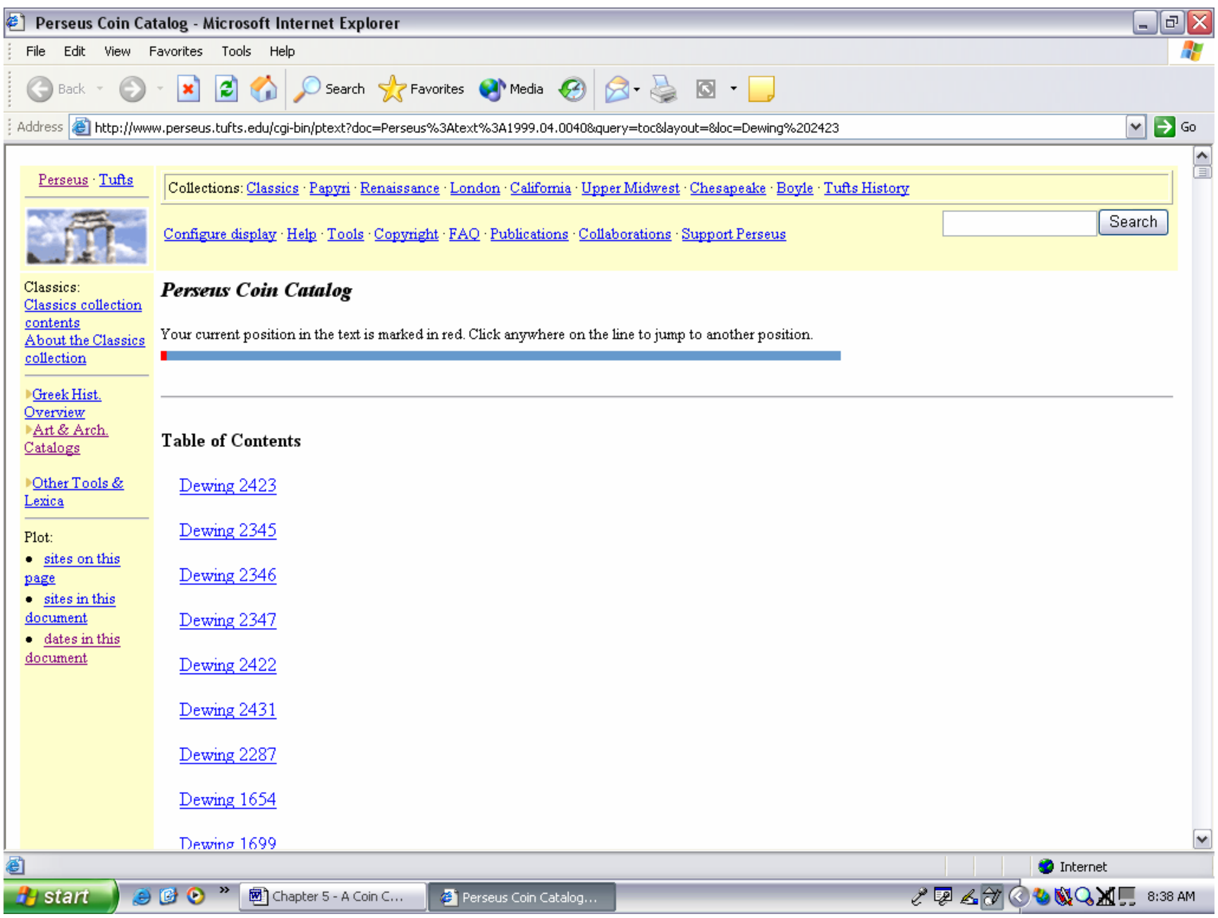

FIGURE 5.1 SAMPLE COIN CATALOG IN PERSEUS

The users can also select a coin from a display of thumbnails, as shown in FIGURE 5.2, which appears to offer a better sense of context in the collection. Unfortunately, thumbnails contain a list of all the images associated with the current table of contents and may or may not reflect the specific coin collection accurately.

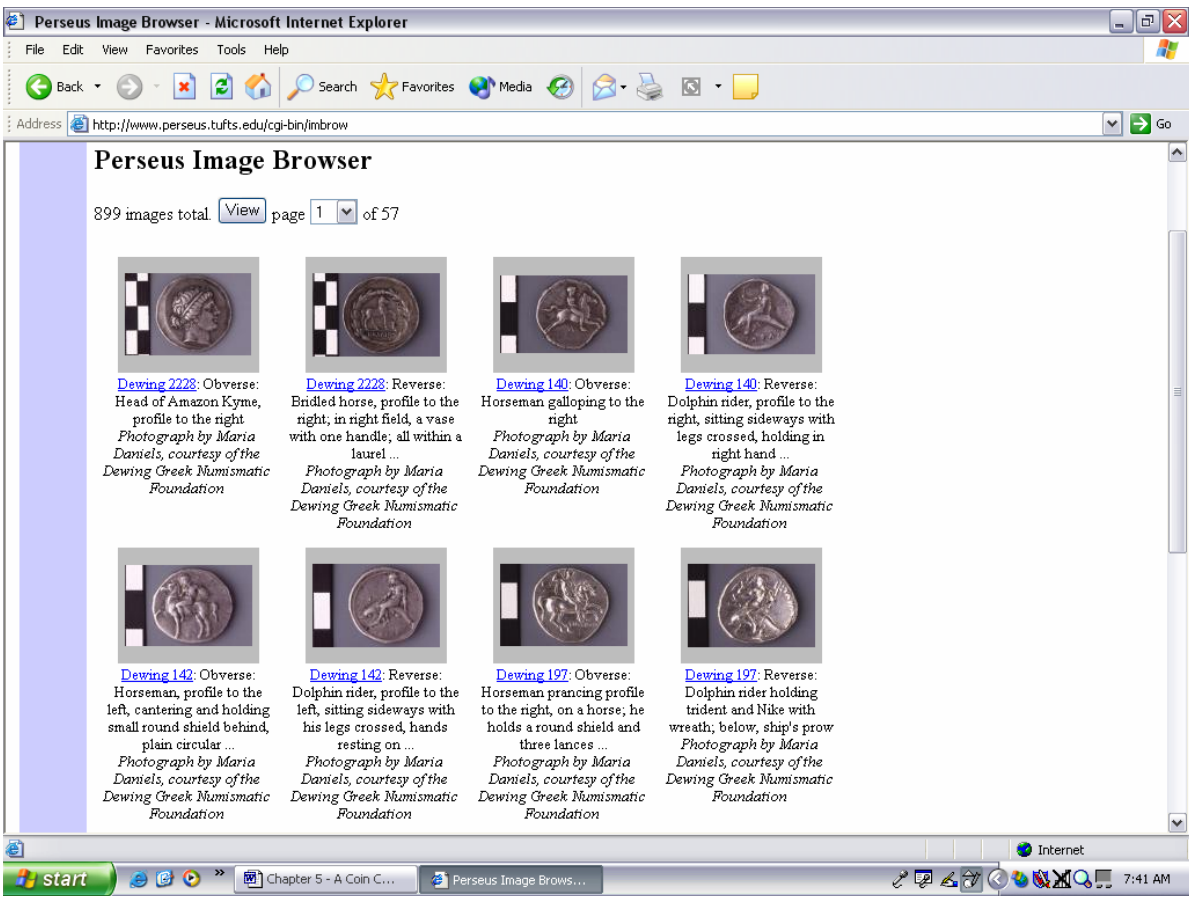

FIGURE 5.2 SAMPLE COIN THUMBNAILS IN PERSEUS

Two images showing the front and back of each coin are generally displayed. A measuring tape is frequently placed next to the images of the coins to give users a sense of dimensionality. Since coins are simple in shape, users can form a good mental model from these two images. The list and thumbnails are the two most frequently encountered approach for displaying coins in the catalog. Users can select any coin of interest and examine it in greater detail. A typical HTML page of the coin shows the collection, material, denomination, mint, date and other information, as shown in FIGURE 5.3.

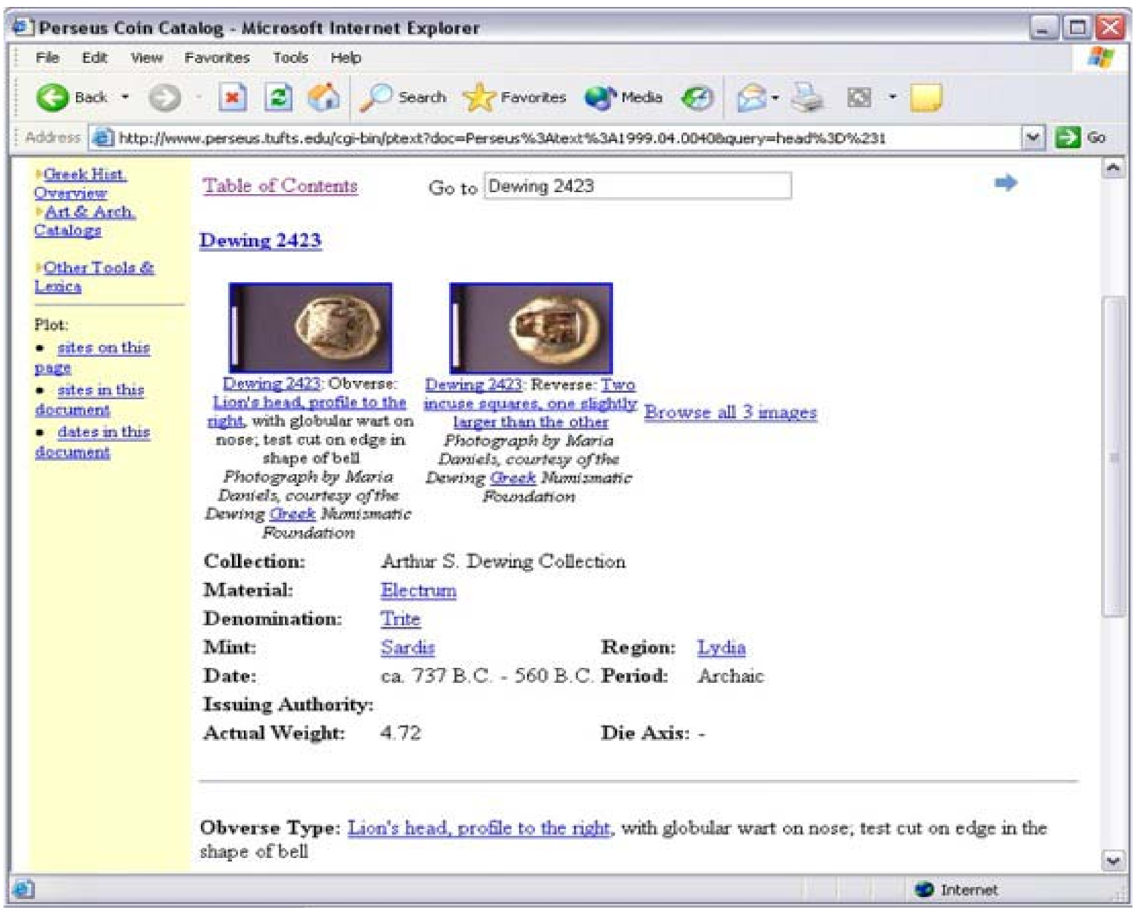

FIGURE 5.3 SAMPLE COIN HTML PAGE IN PERSEUS

The current approach of learning about the Greek coins is very effective in providing focus-based information for a given collection. However, it is ineffective in comparing one or more collections. Users can compare two collections by initiating more than one instance of the web browser. However, since they can only surface one window at a time, the visualization is restricted to the 2D screen space.

5.2 DESIGN RATIONALE

We seek to remedy the deficiency of the current approach and add to its values by providing users a different way of visualizing large data sets (coins) with a common context (collections). The data are homogeneous (all round discs) and supporting data are categorical (years, period, mint, etc.) In light of the nature of the information space, we choose a geometric paradigm to graph the data in selective manipulation in a virtual space. Our design involves a 3D room in the form of a coordinate system. The data points are the data artifacts, and vice versa. Our design concept can be extended to other data domains. For example, the turns of the perspective walls can become the data points for a line graph. We can compare the metrics of the Berlin walls of Germany against the Great Walls of China. Going back to our prototype, we want to come up with a new way of interacting with the data points (the coins) in the virtual environment The 3D information space seeks to provide the users with focus+context spatial learning where they can easily study and compare.

For our prototype, we have modeled 157 Greek coins selected from two collections found in Perseus: the **Dewing** (Dewing Numismatic Foundation) coins and the **BCMA** (Bowdoin College Museum of Art) coins. Both collections contained hundreds of coin. There are a total of 62 Dewing coins and 70 BCMA coins for a total of 132 coins modeled.

The catalog numbers can be found in Appendix B.3. Coins in the collections are modeled from the front and back images found on the corresponding Perseus HTMLs for the coins. The coins are shown floating in a room that simulates a 3D scatter plot. The front and back walls of the room define the distribution of the years (i.e. 700 B.C. – 0). The right and left walls of the room define the periods (i.e. archaic, classical, etc.). It turns out that coins from the two collections are very different. FIGURE 5.4 is a screen capture of the coin catalog in 3D space. The virtual room has the 3D layout by year and period. The ground plane has the 3D layout of materials, such as gold, electrum and silver. It can be quickly determined from the virtual coin catalog that the BCMA collection contained coins made throughout the timelines; whereas most of the Dewing coins modeled are limited to the period between 500 – 400 B.C.

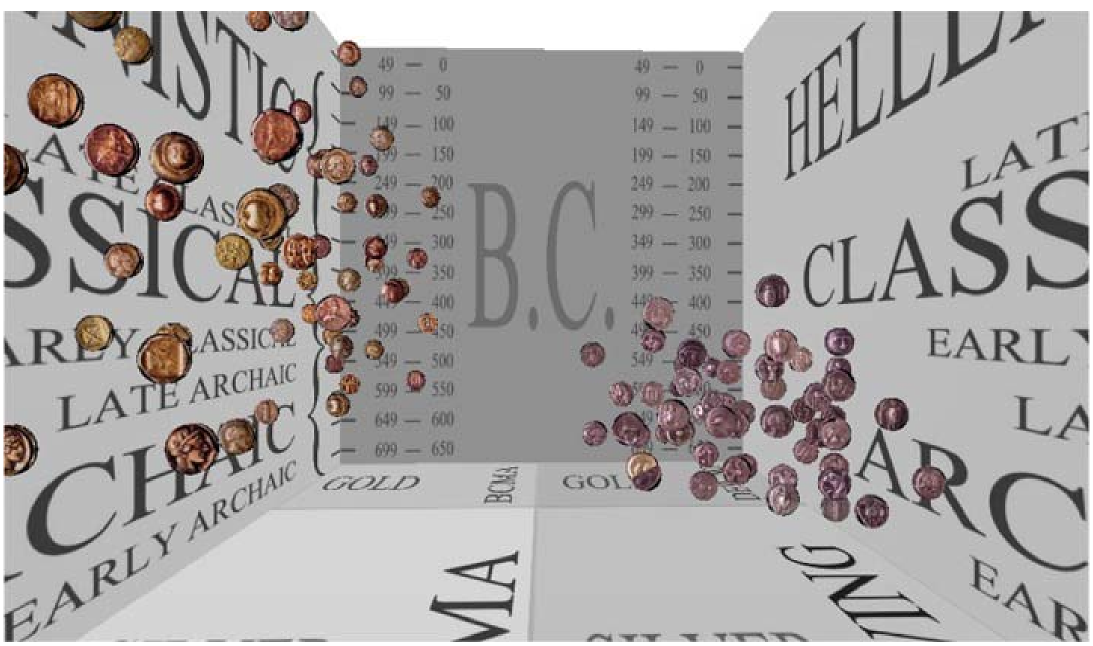

FIGURE 5.4 THE COIN CATALOG

The coin catalog allows users to navigate and explore in the 3D scatter plot. With a HMD, they can be immersed in the graph where the data points are the data artifacts. As they approach a coin, information about the coin can be displayed in 3D space. Using a data glove, they can grab a coin, examine it, and release it. The coin can either snap back to the original location, or it can stay in the user-placed location, in which case, we are allowing the users to re-arrange the coins. In a desktopVR version, with no immersive equipment, the coins can be animated to rotate as users approach or trigger them.

FIGURE 5.5 shows a bird's-eye view of the coin catalog. Although the color tones may appear to suggest that the coins are made from different materials, the users can easily see from the graph that the majority of the Dewing and BCMA coins are silver. There is only one gold coin, which happens to be Dewing. There are also some Dewing coins that are electrum (a gold and silver alloy). All BCMA coins are silver. In a 3D world, the spatial information is self-evident and can be obtained almost instantly. Such contextual

FIGURE 5.5 BIRD'S-EYE VIEW OF THE COIN CATALOG

information cannot be easily determined in a 2D space without more time investment. Currently, the coins are facing the screen. As the users move up in the Z-axis, they will eventually see only the sides of the coin, as already shown in FIGURE 5.5. This problem can be easily resolved by providing a billboard action with proximity sensing; that is, if the users are far away, the best side of a coin will always face them. This is to attract their attention and encourage them to explore in the area of interest. However, once they are within grabbing range, the coins should remain static so that users can 'look around' and examine them in detail. As users grab a coin, other coins in the scene with similar interest can be highlighted with light intensity or rotation animation [107, 109]. Since the color and intensity of the coins should best not be changed, we can use transparency and blur as selective cues for adding layers of coin collections. [50] For example, matching coins with similar nature are sharper in focus while non-relevant coins be more transparent or blurry.

Since the 3D scatter plot can be a very large space, it is very important to allow users to move in all directions. They should be able to conveniently peter-pan in the room, zoom to a coin, or find a coin from a view port. Although only two collections are modeled in our prototype, the coin catalog can be expanded to include more than two collections. We can add or remove data layers represented by different collections. There can be more control in the arrangement and rearrangement of the coins, timelines, and other graph metrics. The design concept of the coin catalog, where the 3D models are the data points, can be extended to other graph formats. For example, we can categorize coins from different collections by a timeline graph. The coins can be sorted by their market values. We can also scatter the coins by geographical regions to show the distributions of currency. We can also mix the coins with other data artifacts. For example, we can show the relationship between coins and the reigns that issued them. An ideal implementation should support both desktopVR and immersiveVR with the arrangement and rearrangement of the collections in real time.

5.3 IMPLEMENTATION ISSUES

The 3D room is made from five thin boxes (the floor and four walls) oriented in an open cube. The walls and floor are texture mapped with an image of the graph layout. In a full-functioning system, the graph layouts should be dynamically generated in real time. The room can also be modeled from a box with the top facet removed. The overhead of a single box is more expensive texture-mapping by facets. The original version of the 3D room is in color (each wall has a different color) and in grids to simulate the XYZ coordinate system. In the final version of our prototype, we have opted for a gray room to reduce color noise and distractions caused by the lines of the coordinate system. We want users to focus their attention on the coin clusters. To that purpose, the background of the room can be further diluted.

Another problem with the 3D room is that users can go beyond the room and into the black abyss of a virtual environment. There are many inexpensive solutions to this problem. We can make the room very large, with a large perimeter for navigation. We can also set boundaries to the navigation domain. We can stop the users from leaving the room by using collision detection or proximity sensing. We can also create a room-like setting using a global background, which is in essence an infinite box.

Intuitively, one would think that they could be modeled from cylinders with a small Z height, which can be made to resemble the thin coin discs quite realistically. Unfortunately, the default surfaces for the cylinders are the column body and not the top and bottom caps the coins need. In order to texture-map to the top and bottom facets, one would have the bypass the defaults and texture-map user-defined facets in a material library. Such effort seems to be an over kill. Instead, oil tanks are used to model the coins because the top and bottom facets are the default texture-map surfaces. The coins are reconstructed by texture mapping photographic images of the coins onto the oil tanks. The coins are placed randomly in a 3D room floating in midair.

In the coin catalog, there are 432 objects with 158,968 vertices and 316,860 faces. The physical memory usage is between 373.3 MB (minimum at frame 0) and 511.5 MB (maximum). The virtual memory usage is between 363.6 MB (minimum at frame 0) to 1,506.3 MB (maximum). It takes approximately 00:00:24 seconds to produce the view shown in FIGURE 5.4 at the maximum frame rate of 100 (see Appendix A, Performance). Navigation is too slow to be practical because the coins are clustered mostly in two large groups visible in one view port. Although the oil tanks are not complex meshes, the sheer volume of the geometries in a single display slows the system to a snail's speed given available computer resources (Appendix A, System Requirement). Navigation is too slow to give a sense of immersion and interactions. To compensate for the limitation in speed, camera views are absolutely required to move from scenes to scenes. Each coin is animated to rotate on its own in 3D space to offer the front and back views.

For the sake of speed, some modeling issues that are glossed over in the coin catalog are the irregularity, side groves, low relief and wear-and-tear. An irregular shaped can be modeled as meshes. It can also be simulated using techniques in negative texture mapping. The sides of the coins can be simulated by shrink-wrap a texture-map JPEG (JPG) file over the coins in the positive normal. The perimeter of the texture-map JPEG file should have a groove-looking design. Low relief, or the slight protrusion of the artwork, can be done by mesh distortion. The wear-and-tear of the coins can be simulated by adding slight bump maps to the exiting texture maps. Realism can also be improved if better pictures of the coins are available. As the number of coins in the scenes increases, a major hurdle in the realization of the design concepts is the slowness in rendering the scenes in real time. Therefore, advanced modeling is not pursued. Instead, some of them are simulated using less costly methods. For example, irregular shapes are simulated by black backgrounds. 3D relieves and grooves are simulated by sharpening and shadowing the texture images.

The original attempt is to implement a 3D scatter graph with support in HMD, a 3D mouse and data gloves. Users can navigate in the 3D space in real time and grab the coins of interest. The system can arrange and re-arrange on command. Unfortunately, with 132 coins, the system crawled to a snail speed. Although it is still possible to view the scene in a HMD and navigate with the arrow keys, the slow real-time response defeats the purpose. To compensate for the deficiency in motion, animation is used to rotate the coins.

5.4 SUMMARY AND FUTURE WORK

In summary, the coin catalog seek to provide added benefits to the existing approach by creating a 3D virtual environment where one or two collections of coins can be easily compared. We proceed to do so by defining the focus of the virtual environment to be the coins and the context to be the collections. We will take advantage of the added screen space in VR and display clusters of coins by collections. The coins are reconstructed true to the nature of the original pictures. Users are allowed to dynamically explore the information space and interact with the Greek coins, which are the data points of 3D scatter plot. The visualization used is 3D geometric and the interaction style is 3D interactive navigation, animation and data glove grabbing. The benefit we seek to add is to allow users to compare one or more coin collections through focus+context interactive visualization in a virtual environment. A coin is the smallest focus and the 3D scatter plot is the context. Each collection is the larger focus or the intermediate context. Users are able to maintain focus and different levels of context simultaneously.

More than be done for the coin catalog besides the cosmetic improvement already covered in the implementation issues. Currently the 3D scatter plot is merely a layout with no loading algorithm. Future versions should take into consideration self-organization, automatic loading and advanced plotting techniques. [14, 17] We can also superimpose more than one graph style (i.e. line graph on scatter plot to show trends) in multi-dimensional layouts. [64]. Future work should allow users to go into the detail of each coin within the context of the graph [108], perhaps shifting from 3D to 2D paradigms, or other advanced techniques such as world-within-world [65].

Current performance can be greatly improved with advanced implementation techniques. With more programming investment, there is no reason why it should not be efficient enough to support both desktopVR and immersiveVR. One of the greatest benefits of homogeneous datasets is the congruency in data entities, which allows efficient data management [87, 88] of the texture maps. The coin catalog is also suitable for object-oriented programming where instances of the same coin module can be used to populate a large scene. The major challenge appears to be in the arrangement and rearrangement in real time. Since there must be enough clearance between the coins to allow comfortable navigation and the design is grid-based, perhaps refreshing speed can be enhanced by applying some concepts in raster operation. That is, the maximum number of coins supportable by the scene is preloaded and hidden. Based on the user requests, coins in the appropriate location can be 'turned on' with the proper texture maps. For small-scale or locale refreshing, moving the coins in virtual environment can be the most efficient and least costly approach.

CHAPTER 6

3D INTERACTIVE GREEK VASE MUSEUM

The 3D interactive Greek vase museum (vase museum) is a virtual reconstruction of selected ancient Greek vases in the London collection in Perseus. In this chapter, we will present the design rationale and implementation issues in the modeling and renditions of the vase museum. In the introduction section, we will closely examine the focus+context problem as applied to the Greek vases. We will present a comparison of our new approach versus the current approach. The current approach of learning the vases is by browsing the HTML pages and reading the text description in Perseus. The new approach of learning the same collection of vases is by exploring a 3D interactive museum. We have created a high-fidelity prototype of our design concepts to demonstrate the suggested advantages. In proof of theory, we have conducted an experiment in desktopVR to evaluate the improvement and to examine some of the elements that account for user performance with the new focus+context interactive interface, by comparing our prototype to corresponding conditions in Perseus. In the next chapter, we will present the evaluation of the vase museum and the results of the usability test conducted in comparison of Perseus.

6.1 INTRODUCTION

The current approach is epitomic of the focus+context problem for several reasons. First, the Greek vases are beyond the reach of the museum visitors and must be viewed from a safe distance. Therefore, the digital library serves as an important alternative source of information and of visualizing data artifacts. [16] Second, there are various levels of vase data quality in the digital library. Most have one or two photographic images. Some have no photographic images; some can have as many as 31, as shown in FIGURE 6.1.

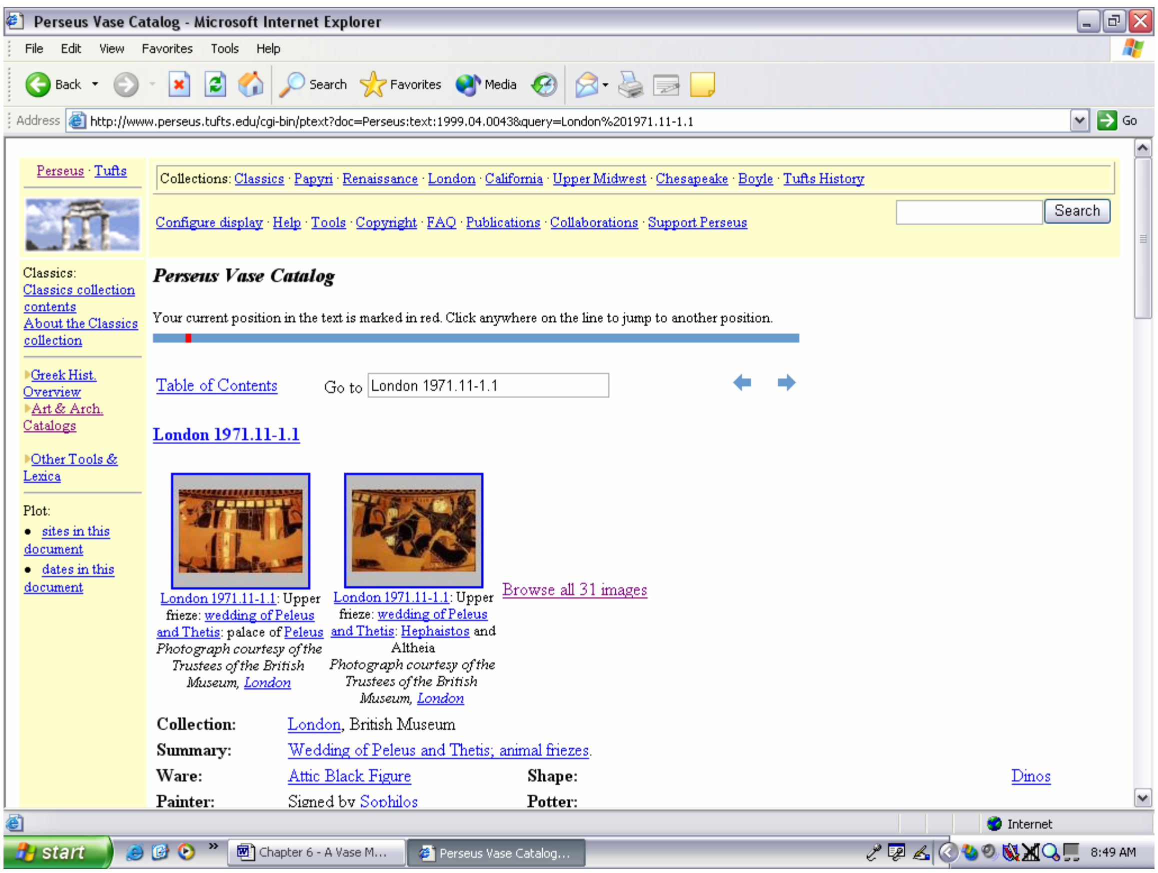

FIGURE 6.1 SAMPLE GREEK VASE IN PERSEUS

Third, vase data can be severely crippled in the digital library. Not only are the photographic images piecemeal and fragmentary, the users must rely on awkward text description in the absence of vital visual information. Long paragraphs are used to describe the appearances of the vases, such as the number of handles, openings, description of artwork, and so on, which the users can obtain in one glance at the real vases. Forth, there is no effective method of maintaining focus (vase) and context (collection) on a 2D screen space. In accordance with current convention, the users click select from a list of vases. FIGURE 6.2 shows a screen capture of a list of the London collection in Perseus.

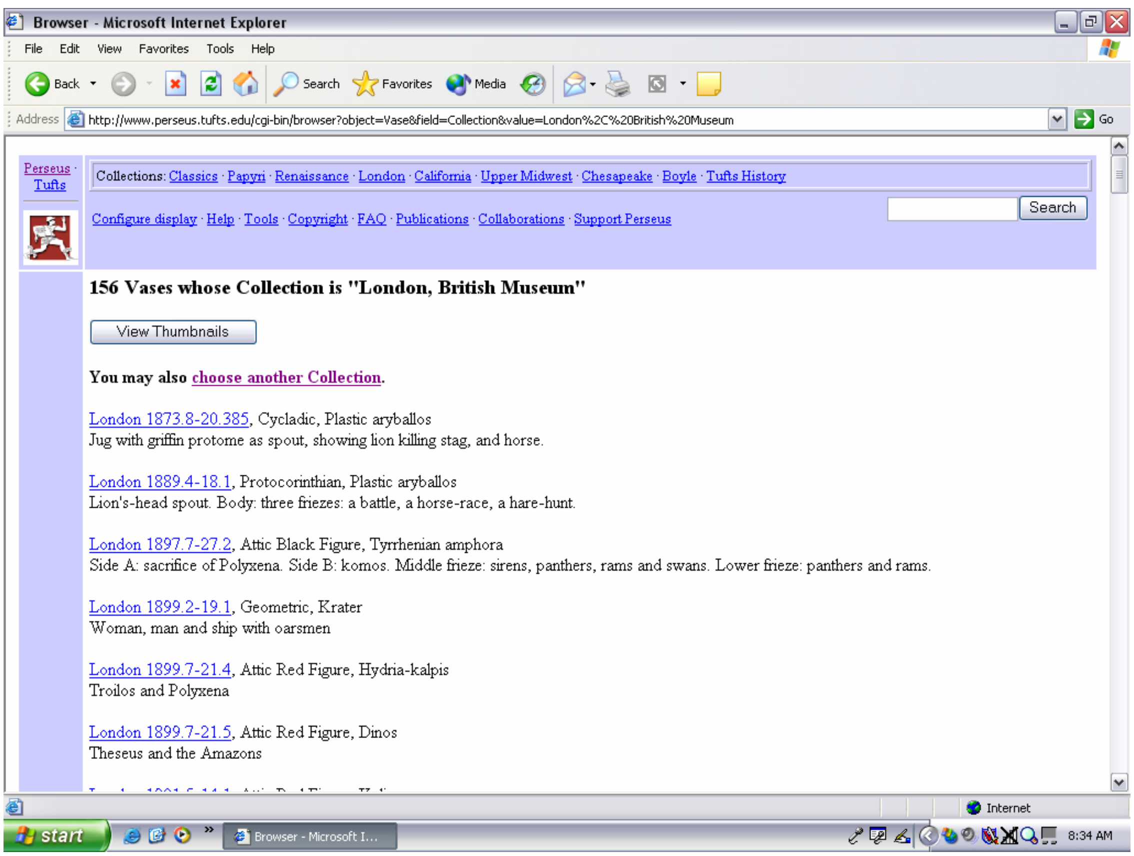

FIGURE 6.2 LONDON COLLECTION OF GREEK VASES IN PERSEUS

It may seem that a thumbnail view could be a better approach to maintain focus and context in the vase museum, serving almost like a zoomable browser. [2, 93, 154] FIGURE 6.3 shows a thumbnail view of FIGURE 6.2. Unfortunately, the thumbnail view shows all 221 images associated with the 156 vases. Vases with multiple images appear to be representative of the collection and vases with no images disappear.

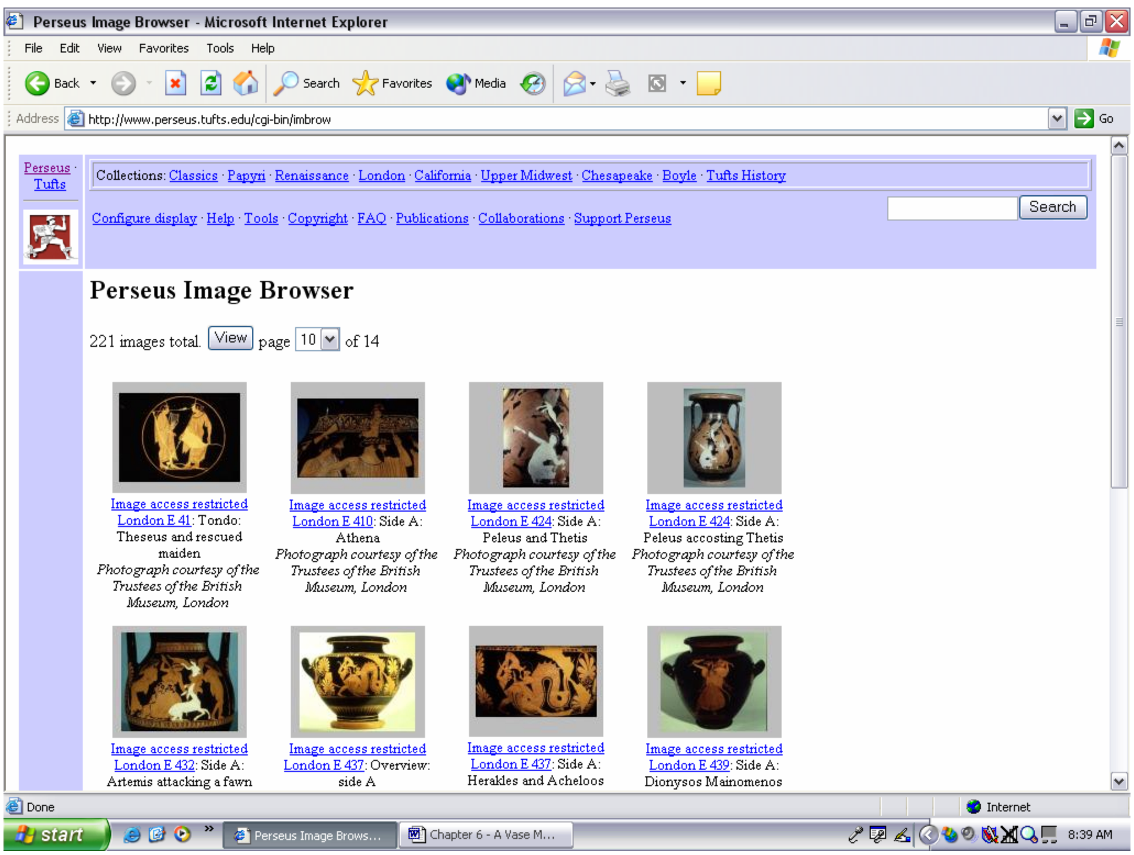

FIGURE 6.3 THUMBNAILS OF LONDON COLLECTION IN PERSEUS

Finally, the users cannot correctly assume the general appearance of the entire collection by sampling a few vases. Reading the text descriptions may not be as helpful to untrained novices because archaeological terminology can be prohibitive. Unlike other data artifacts (buildings, maps, coins, gems, sculptures, etc.), which the users tend to have a preconceived notion of general appearance, new users to the Greek vases must rely on the photographic images of the vases taken at different angles to get a mental picture of the Greek vases.

6.2 DESIGN RATIONALE

We seek to remedy the current approach by using virtual 3D as a learning environment. We will provide the users with an interactive museum where 3D models of the vases are on display. The users can walk about, look at and 'touch' the vases. The benefit we seek to add is 3D focus+context visualization of the vase artifacts. A vase is the smallest focus and the museum defines the context of the entire collection. As the users walk about in the museum and explore the vase displays, they can acquire spatial input that may not be as efficiently provided in HTML pages describing the vases. The context of the collection is provided by the information landscape where all the vases in a collection can be arranged and rearranged on user command. While the users focus on one vase, they will not lose context and relational reference with the other vase in the scene.

For our prototype, there are 157 vases modeled in the London Collection. Appendix B.4 contains the list of the catalog numbers of the vases modeled. At the time of the prototype modeling, there are 156 vases listed on the site, but in reality, there could be more.

Since Perseus is an evolving project, the sites are frequently modified and the catalog numbers may not be an exact match to the current display. In order to have a realistic museum, it is a must that 3D models of the vases are used to populate the museum. By restoring the 3D nature of the data artifact, we can eliminate the need to describe the appearance of the vases in text. The 3D models add value to the existing system. Currently, when the users click on the image of a vase, they can at best get an enlarged version of the picture. In comparison, our prototype system can show a 3D model of the vase image. As shown in FIGURE 6.4, the users can move and rotate the 3D vase to study it.

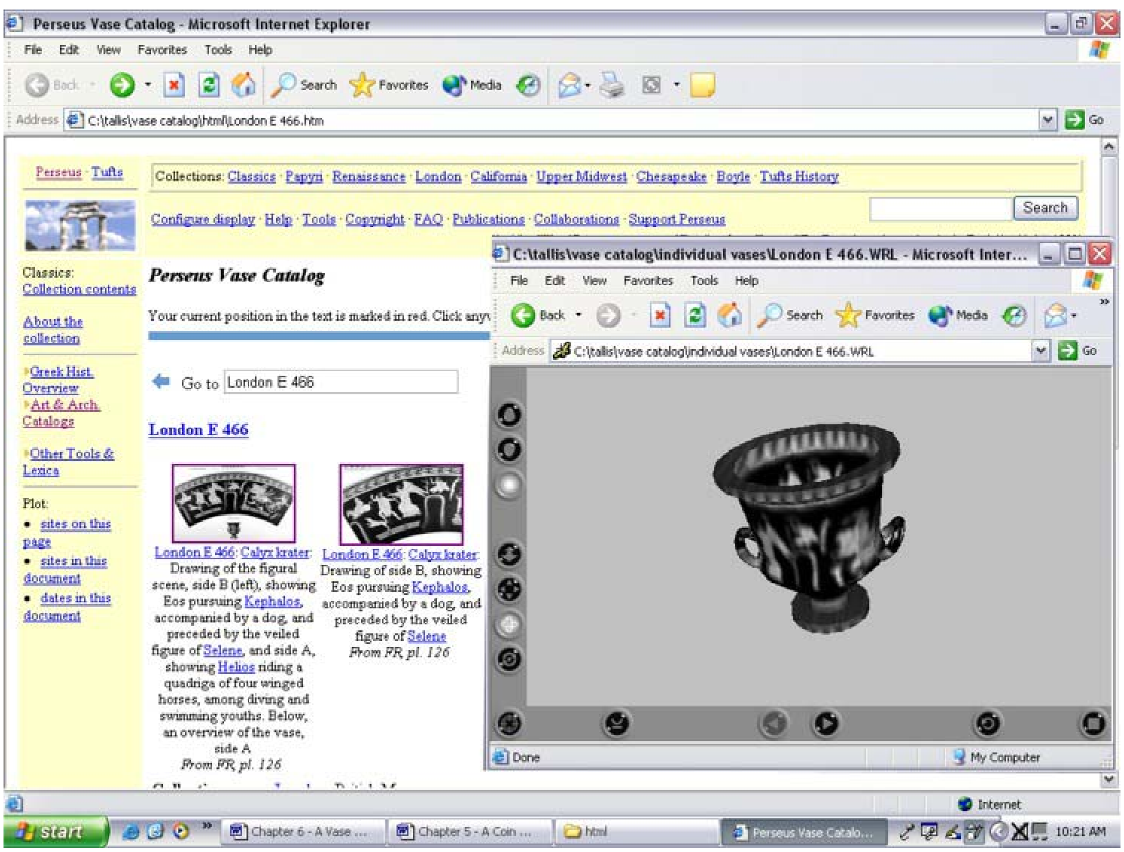

FIGURE 6.4 HTML PAGE AND 3D MODEL OF LONDON E 466

We have built two renditions of the vase museum. The first rendition, shown in FIGURE 6.5, is a crude simulation of a real museum, complete with texture-mapped floors, ceilings, walls and exits. A random selection of vases is displayed on marble columns with captions. The users can navigate in the museum and browse the vases. The simulation is similar to existing approaches in dynamic and interactive explorations. [5, 200, 202] Ideally, the users can pick up a vase, examine it and put it back on the display column. A museum curator should provide user help. The users should be able to arrange the room, sort the vases or group the vases in certain categories. The wall spaces could be used as secondary information space. An ideal application of a realistic vase museum should encompass information on the museum settings itself as well as the data artifacts. For example, we can model the British Museum where the London collection is on display.

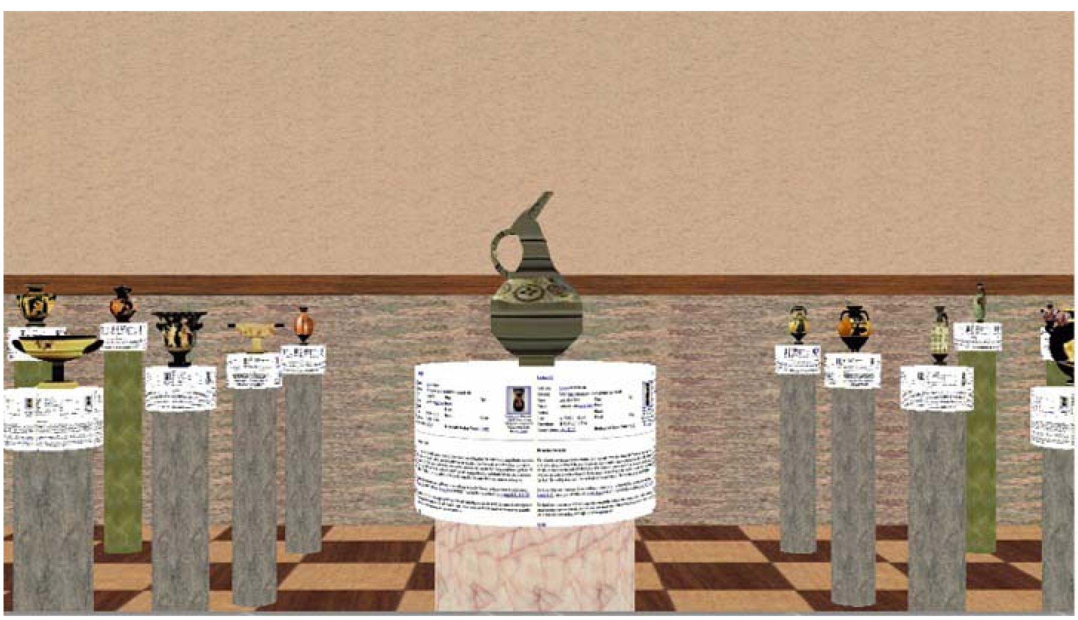

FIGURE 6.5 REALISTIC RENDITION OF THE VASE MUSEUM

In the second rendition of the vase museum, shown in FIGURE 6.6, we depart from the conventional realistic approach and seek to build an environment that can bring out the focus in a more meaningful context. The background is grey-scaled to reduce as much background noise as possible so that user focus can be directed on the vase artifacts. [90, 129, 130, 135] The vases are placed on very thin columns that are only slightly tinted. The tints of the columns represent different source regions. Unknown regions have grey columns. Vases with no available images are shown in grey models. The walls show year on one side and ware (type of glazing) on the other. The vases are placed in a year by ware layout. In doing so, the vases form a natural clusters by the type of glazing made for the specific year periods. Ideally, the users can arrange the vases on command in real time and visualize the vase in a different geometric layout by various parameters.

FIGURE 6.6 ABSTRACT RENDITION OF THE VASE MUSEUM

The vases in the London collection mostly fall into two main types of ware (both with subtypes): the red figures and the black figures. FIGURE 6.7 shows a perspective view of some red figures in the vase museum. FIGURE 6.8 shows a perspective view of some black figures in the vase museum.

FIGURE 6.7 SOME RED FIGURES IN THE VASE MUSEUM

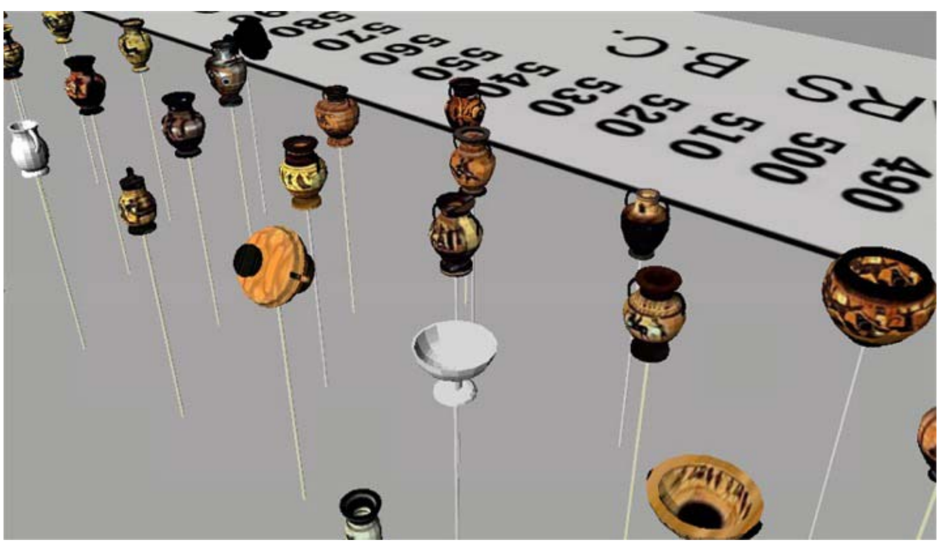

FIGURE 6.8 SOME BLACK FIGURES IN THE VASE MUSEUM

Looking down from the ceiling, the placements of the vases reveal a 2D scatter plot by year and ware. FIGURE 6.9 shows a bird's-eye view of the vase museum looking straight down. It shows the two main types of wares made throughout 700 – 300 B.C.

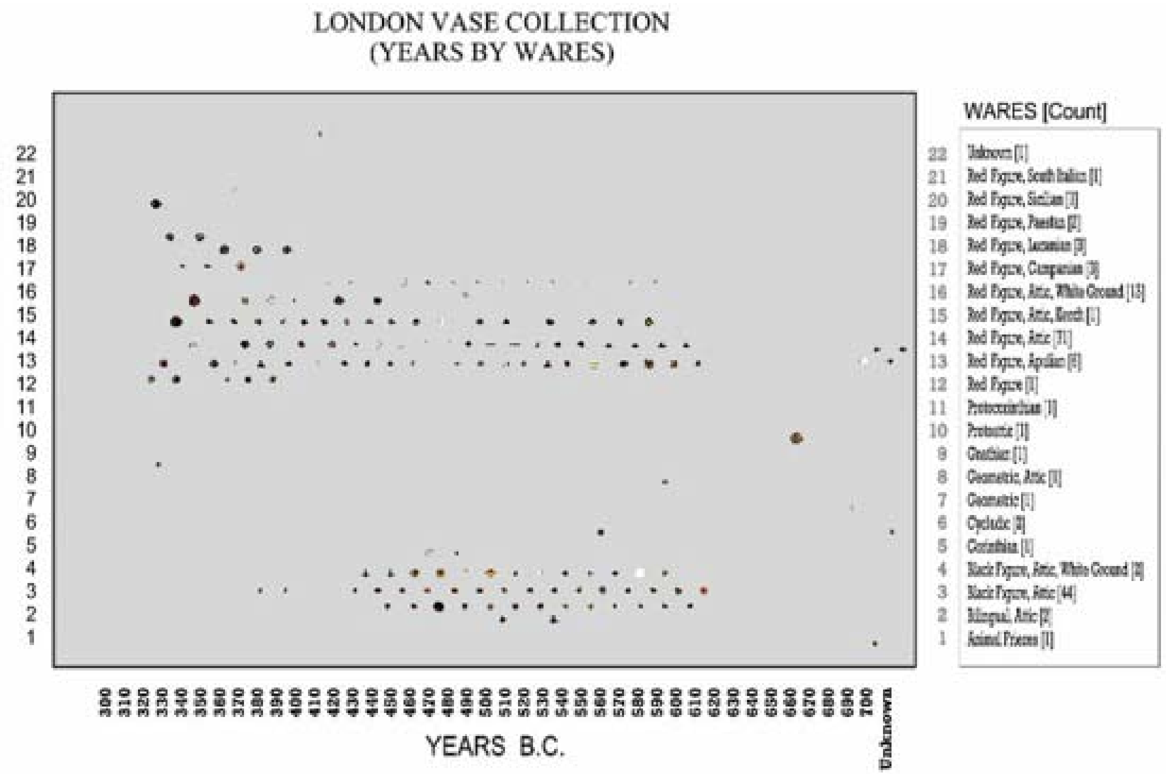

FIGURE 6.9 3D SCATTER PLOT OF THE VASE MUSEUM

Proximity sensors in the vases can determine location of navigation. Close proximity implies user interest. When the users are within a certain distance, information about the vase appears in the scene. More than one proximity sensor can be triggered if the proximity perimeters overlap. If the proximity perimeter is set to be very small, only one vase can be triggered at any given time. For demonstrational purpose, the information displayed for the prototypes are the corresponding HTML pages, although other types of information could be similarly displayed. For example, we could have a world-in-world metaphor where the users can trigger and go to a different scene telling a story about the vase. FIGURE 6.10 shows the triggered proximity sensor for vase B 376.

FIGURE 6.10 LONDON B 376 IN THE VASE MUSEUM

As the users approach the vase to get a closer look, a semi-transparent HTML page appears on the right. A semi-transparent text caption shows up on the left. As they approach closer to the HTML page, it becomes less and less blurry. [50, 127] Each HTML pages contains an anchor that can jump to another scene. If the users click on (in desktopVR) or touch (in IVE) the HTML page, a shift of paradigm occurs and the system loads the 2D version of the HTML page, with full support of the existing the digital library techniques. The users can read about the vase with the interactive museum in the background. When they are done, they can return to the interactive museum and continue to navigate in the room. When they move away from this vase, the associated information disappears and other proximity sensors can be triggered. The users can turn off the proximity sensors. In that case, they can walk about in the room without transient information. To ensure that the users always get a clear view of the text and pages, billboards are implemented so that they will always turn around the vases facing the direction of the users. The vases in the room form the primary information space. The 3D HTML pages form the secondary information space that is superimposed on the primary. Leaving the virtual scene and returning temporarily to the 2D HTML pages forms the tertiary information space. Since some HTML pages contain more information than others do, as the users navigate the interactive museum, the transparent pages may vary in size, which implicitly informs the users on how much we know about the vases. This piece of secondary information cannot be obtained in Perseus or existing other digital libraries.

6.3 IMPLEMENTATION ISSUES

Our prototypes put a heavy demand on existing computer graphics capability and rendering speed. The vases are modeled individually from 2D pictures and stored as VRML files using the same name as the vase catalog numbers. The 3D models are inserted into the vase museum by using the X-Reference feature in the modeling program (Appendix A). All related HTML files are stored in one directory. The texture-map files are stored in another directory. A full-function, real-time system for much of the conceptual designs of this thesis is not yet achievable given currently accessible resources.

In the vase museum, there are 2,155 objects with 260,411 vertices and 510,614 faces. The physical memory usage is between 464.2 MB (minimum at frame 0) and 511.5 MB (maximum). The virtual memory usage is between 417.2 MB (minimum at frame 0) to 1,506.3 MB (maximum). Given the computer resources available, the vases cannot be displayed in a single view due to insufficient virtual memory (Appendix A, System Requirement & Performance). The users are fixed in a walking mode, with a speed of 2.0 and a height of 75 units per second. The museum room is set to be relatively large in a scarcely scatter plot, such that the number of vases displayed in any given view port is limited as the users navigates the room. Navigation in a large room can be slow and ineffective. The users may become impatient in a long walk. The following sub-sections will discuss some of the implementation issues in the making of the vase museum.

6.3.1 VASE MODELING

The vases are reconstructed as complex meshes true to the nature of the original photographs in Perseus. The vases are constructed from geometric lathes of 2D line drawings of the photographs found on corresponding Perseus HTMLs. The lathed models are texture-mapped with a JPEG image file of the vase created from the same photographs.

The London collection can include vases, lids, plates, cups, and a composition, as shown in FIGURE 6.11.

FIGURE 6.11 LONDON VASE F 90

Since there is no front or back to the vases, most pictures on the vases show the 'side A' view and/or the 'side B' view. Some pictures show the exterior or interior views (i.e. cups and bowls). Other pictures show the top and bottom views. There is hardly any picture of the side view or of the handles. The lack of vital information on the vases translates to arbitrary and artistic interpretations of the vases in the making of the 3D models. In our prototype, the vases are modeled by side A only in U=2 and V=1 shrink wrap. The UV texture-map setting repeats the pattern for front and back around the body of the vase. To cover the top and the bottom, extra margins are created in the JPEG (JPG) files to cover the brim and the base. Theoretically, it is possible to draw half of the vase and lathe it around the Y-axis to take advantage of the symmetry. In practice, the simplicity will result in vases with positive normal and no interior texture mapping. If the entire exterior contour of the vase is drawn, including the thickness of the brim, the interior will be a mirror of the exterior. The openings of the vases should not be drawn, so that the lathing could produce a hollow mesh (as compared to a solid and filled geometry).

The handles are made from bent toruses or cylinders. Most vase handles are symmetrical. Therefore, only the left handle is modeled because it is easier to draw the left side using the right hand. The right handle is a clone of the left handle flipped or rotated. FIGURE 6.12 shows London E 157, which has the most complex handles in the vase museum. They are constructed from four sets of modified meshes (two for each side merged).

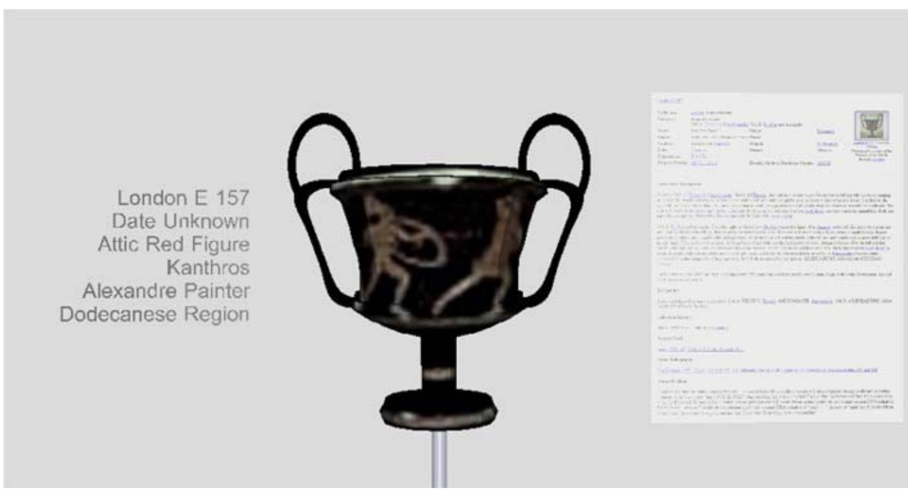

FIGURE 6.12 LONDON VASE E 157

Hiding the handles can be a challenge because the vases are 'thin' and the toruses or the cylinders placed at an angle may protrude and show in the inside. To circumvent this problem, the cheapest solution is to thicken the walls. If that fails to do the trick, then the toruses or cylinders must be converted to meshes, and thus increasing the cost of the prototype. Triangulation of the meshes allow the protrusions to be deleted, which may leave 'holes' if not done tactfully, or 'tucked in' by pulling the vertices of the offending triangles into the thicken walls of the vases.

X-references (Appendix A) are used to merge the vases into the various renditions of the vase museum. All vases are modeled in a source file at origin (0,0,0). X-references allow the changes to the source vases to be propagated automatically in the referenced scenes. X-references seem to work well in most cases. However, a few vases with complex mesh handles, such as London E 157 in FIGURE 6.12, appeared to be distorted and/or dislocated, even though such defects do not exist in the source file. The problem is probably caused by mesh manipulation implemented to preventing the handles from protruding from the interior of the vase. The only solution to this problem seems to be merging real copies of the offending vases into the target scene, losing the advantage of the X-references. Since there are only a few vases with damaged handles, no remedy is taken.

6.3.2 MUSEUM LAYOUT

In the first version of the prototype, the museum room is modeled from five thin boxes (the floor and four walls) oriented in an open cube, and the columns are cylinders. There are aisles between the columns to allow natural navigation. The users have a natural tendency to avoid collision with the columns or the walls. The first version contains only a subset of the 157 vases. When the room becomes large, the context can be vague without labeled regions. It is also difficult to cluster vases in a meaningful way. In the second version of the prototype, the museum is on a floating plane, texture-mapped with the coordinate system of the graph. The walls are backgrounds that are technically a very large box, texture-mapped with the labels of the years and wares. Attempts to render the box hugging the floating plane are unproductive; therefore, it is possible for the users to fall off the plane. However, the margin is large enough such that it never happened in the pilot runs or during the usability test. Another problem is the perspective distortion of the walls with vases. For example, vases lined up with the 500 B.C. regions can appeared to belong in a different a period if viewed from far left or far right

6.3.3 INTERACTIONS

Interactions in immersiveVR are not as fully implemented as in the desktopVR. Viewing the vases in an immersiveVR can be a serious disappointment due to the low resolution of the current affordable HMDs. Proximity and anchor sensors are mouse events and do not have direct translation in immersiveVR with no keyboard. To achieve the same effect available in desktopVR, proximity sensing in immersiveVR must be implemented as part of the collision detection algorithm. Collision detection is also used to detect an intersection of the data glove location and the vase location. Grabbing is equivalent to rotating the vases based on hand gesturing. In desktopVR, rotation of a single vase is easily achieved in shifting the paradigm from 3D to 2D. Rotation of a large scene is prohibitive due to system speed and existing browser limitation. Currently, it is not possible to rotate only a vase in a large scene without major programming investment. To compensate for this deficiency, the vase museum allows the users to switch to the 2D HTML mode. When they click on the vase image, the corresponding vase model is loaded. The users can rotate the single model in very good speed while the vase museum is still running in the original window. The prototype does not allow rotation or peter-pan movements. Theses features are deliberately blocked to avoid what appears to be system freezing. Camera view ports are created for each vase, which allow the users to change the scene to the location of a selected vase quickly.

Navigation can be slow and awkward. During the pilot runs and usability test, it is found that the users tend to click on a vase in an attempt to zoom to it. Distance seems to be interpreted differently in a virtual 3D space. The users may not want to 'walk' to a vase they see in smaller perspective (further away). They may believe that they are 'already there' and transfer existing habits from the desktop metaphor. To allow the users the click-zoom interaction, we could place an anchor in each vase. When the users click on the vase, it would change the camera view port to the location of the selected vase. The system is quick in switching the view port, but the effect can be abrupt and unrealistic. Moving in the X-axis (right and left keys) and Y-axis (up and down keys) are more intuitive than moving in the Z-axis. Moving in all three axes is a challenge. Navigation is incremental in the direction of the movement. A small increment can prolong the walking experience unnecessarily. A large increment can jump over a group of closely clustered vases. The current increment setting jumps over a single vase to prevent the users to walk into the interior of a vase. Moving too closely to a vase can cause the images of the vases to appear pixilated.

There is an animated view port, called **anitour,** which can take the users on a guided tour in the interactive museum. Current animated tour is quick and jumpy because a large distance is being covered in just a few frames. The use of animated tour is debatable in an immersiveVR because it can defeat the purpose of interactive exploration. The users are encouraged to move about dynamically toward interesting vases without any system guidance.

6.4 SUMMARY AND FUTURE WORK

By providing a 3D vase museum with lightweight interactions, we seek to blend some of the benefits of human-computer interaction with those we can get from computer graphics and virtual reality. We seek to preserve the existing functionalities of 2D image viewing, while adding advantageous features that virtual reality affords. We allow the users to focus on a particular vase of interest without losing context of the entire collection. We allow different views of the vases with proximity sensing and easy shifts between the 2D and 3D paradigms. The visualization is 3D data landscape and the interaction styles implemented is 3D navigation with lightweight interactions. New visualization of the vase museum is the 2D scatter plot, which is manifested by placement of the vases. The goal is to provide focus+context view of a selected scene in an interactive museum and allow the users to learn by immersion. Organization of the vase is inherent in the real-life museum metaphor and geometric representation by year made and wares. The users can learn about the vases by interactive exploration.

The vase museum can benefit from improvements in the sense of realism and immersion. Realism of the 3D vase models can be improved by high-resolution images, accurate measurements and advanced modeling techniques. The vertices of the meshes could be stored in a flat file or even in a database. A program should be able to automatically load and populate the virtual scene. Continued effort in the area of immersiveVR can allow the users to navigate more naturally with a 3D joystick or a 3D mouse. The users should be able to grab a vase with a data glove. Ideally, the users should be able to click on a vase in the plot and peter-pan into the scene. The users should be also able to change the parameters of the graph style, legends, axes and other chart information. This is not technically feasible given the current speed of the system in use and the number of vases we wish to have in the scene. However, it is possible to achieve scripted changes by switching from one pre-made scene to another. In practice, that effect is not very satisfactory because the long loading time destroys the sense of immersion. In addition, the system in use did not permit loading two such large virtual environments. Given current computer speed, parallel processing could be the only effective solution to speed up the system enough for large arrangement and rearrangement.

CHAPTER 7

EVALUATION OF GREEK VASE MUSEUM

Visualization of archaeological data was traditionally achieved by visiting the sites of preservation, such as a museum. However, it was not always possible to go to the physical locations where the real vases were kept. With the advent of the digital library, the problem in accessibility was mitigated by viewing photographic images and relevant documentation on the web. A disadvantage of existing digital libraries was the loss of context. As the users focus on one HTML page about a vase, which occupies the bulk of the screen space, they have trouble maintaining the context of the entire collection. Some users try to compensate for this loss by opening and balancing multiple windows. Others tried to search and sort in a way as to reflect some kind of context. In the vase museum, we seek to compensate for this loss by allowing the users to walk about and focus on a particular vase of interest without losing context of the entire collection in the background. We believe that interaction techniques in 3D focus+context visualization should offer the users the spatial learning of a context-driven world. To validate our approach, we have designed an experiment to evaluate the usability of the vase museum by comparing it to the corresponding Perseus site on the same collection of vases. In this chapter, we will discuss the experiment, experiment design, subjects, task, procedures, results and a discussion of the results. The design and implementation issues of the vase museum can be found in Chapter 6.

7.1 EXPERIMENT

The experiment attempted to quantify the costs and benefits of the 3D interactive interactions using a spatial metaphor, as compared to the existing 2D direction manipulation using the desktop metaphor. Because a virtual environment can afford the sense of presence by allowing the users to navigate and explore an information landscape, intuitively, we expect the vase museum to provide some benefits compared to the existing Perseus, which has an interface design representative of existing digital libraries. We expect the subjects to perform a context-based task faster than a focus-based task and with greater accuracy, speed and satisfaction. We have designed an experiment in desktopVR and conducted corresponding usability test to measure such benefits in terms of scores and timing.

7.2 EXPERIMENTAL DESIGN

The experiment was conducted in an HCI lab setting, between subjects, and for about 1 – 1½ hour. There were two groups of subjects. One group (the V group) used the vase museum prototype on a laptop PC running the Cortona VRML browser (Appendix A). The other group (the P group) used Perseus on a desktop PC running the Internet Browser. Both groups performed an equivalent task and answered the same set of questions, quiz and surveys. The task involved questions that subjects can answer by visualizing the vases on a computer display. Our plan was to design a list of questions, surveys, and quiz that was equivalent and unbiased to the system in use. [183, 185] A complete listing of the data collection instruments, which both groups used, can be found in Appendix C. The test results can be found in Appendix D.

In our experiment, we asked the users to 'look at' the vases in the London collection as a whole (context). They were to observe similarities or differences among the vases. The users were asked to focus on a particular vase or a group of vases and compare it with the surrounding vases in context. In Perseus, the context was provided for by a list of the London vases. The users could search the list or search in Perseus. They could focus on a vase by clicking on the corresponding HTML page. In the vase museum, the context was provided for by a 3D room. The users could focus on a particular vase by 'walking' to it. Walking could also be simulated by selecting a camera view.

7.3 SUBJECTS

There were a total of 20 subjects divided into the V group and P group. Each group contained 10 subjects, 6 male and 4 female. The subjects were not required to have any prior knowledge of the desktopVR, Perseus, Greek vases or archaeology. However, they must be experienced computer users since timing in performing the task was critical. The distribution of the subjects was maintained in a round-robin fashion. Most subjects appeared to be in their 20's. The subjects were recruited from the department of computer science and the department of electrical and computer engineering at Tufts University. They were solicited by distributing printed ads, posting notices on course web sites, inviting known students, and accosting unknown students. The subjects were not paid, but they were offered a choice of gift trinkets purchased at $3–5 a piece. Two subjects declined to pick a gift.

7.4 TASK

In the vase museum, the V subjects performed the task of "exploring" the 3D room where 3D models of the vases were on display. In Perseus, the P subjects performed the task of "browsing" from a list of all the vases in the collection. While exploring or browsing, the subjects answered the same set of questions, surveys and questionnaire. Subjects were tested for their accuracy and speed. Subjects were instructed not to study each vase in details; instead, they should just get a general idea of the whole collection. They were to concentrate on the appearance of the vases, such as colors, shapes, handles, decorations and artwork. There was no counting, reading, searching or sorting required. Certain actions that diverted the subjects from the task, such as leaving the Perseus or going out of the specific vase collection, were checked and corrected. The task contained 10 multiple-choice questions which were analogous to home work assignments for archaeological courses found on the WWW. Each question was designed to reflect a common characteristic of the entire collection. For example, there was a categorical question for each of the major traits on colors, shapes, on number of handles, etc. The answers reflected the majority context of the collection. Each task question had its own set of 8 associated survey questions on a Likert scale of 1 (strongly agree) – 7 (strongly disagree), so the lower the numbers, the better the response. For each subject, there were a total of 80 survey questions that were averaged. To measure the depth of learning, a survey quiz was given at the end of the experiment.

7.5 PROCEDURE

Subjects were given a very brief introduction to Perseus and the vase collection. They were given a warm-up question as a practice that did not count in the final score. Similar to the actual task question, the warm-up question contained a multiple-choice question and eight survey questions associated with that specific question. The warm-up question was focus-based and asked the subjects to directly select two vases and compare them. In the vase museum, this was achieved by simply choosing the view ports for the vases. In Perseus, this was achieved by clicking on the links from the list of collections. Warm-up time, including system training, lasted approximately 5–10 minutes. In the actual task, subjects were given one question at a time. They were timed while they performed the task. When they were done with the question, the timing stopped, and they were then given a survey of eight questions that pertained specifically to that question, which was not timed. They were allowed to change their answers until the clock was cleared, which occurred during the period while they were performing the survey. There was no backtracking of the questions. Subjects were encouraged to pick the best answer possible, so no task question was left unanswered. User observations were noted.

After the subjects had completed the task, they were asked to complete a questionnaire that had three parts. Part A was a survey of 20 questions based on the Likert scale of 1 (strongly agree) – 7 (strongly disagree), which measured the overall satisfaction of the experiment in general. Part B was a written section for comments about the experiment, such as navigation style, answer approach, stronge/weak points, likes/dislikes, and recommended improvements. Part C was a quiz about the vases. It was placed at the very end to allow the maximum buffer time between performing the task and taking the quiz.

7.6 RESULTS

User observations can be found in Appendix D.1 and user comments can be found in Appendix D.4.2. There was no special trend in subject opinion or system preference. Both groups were able to complete the task. A couple of subjects skipped some questions on the quiz and resulted in bad performance. User observation revealed that the V subjects primarily used the view ports to select a few scenes randomly to observe the vases in the background. Some V subjects 'walked' in the room, many did not even bother. Those who tried to explore the room complained that navigation was too slow to be practical. A few tried to walk to the back of the vases and found that the maneuver in desktopVR using cursors was cumbersome. Only a couple of subjects shifted to the 2D HTML pages. Most subjects ignored the scatter plot in the vase museum. Nearly all complained in the warm-up that the view ports for the vases were not sorted alphabetically. However, since the warm-up question was the only one that was focus-based, requiring them to go to specific vases, subjects did not bring up the issue during the actual testing. A few subjects complained about the touch pad on the laptop and they would prefer a real mouse.

User observation revealed that the P subjects primarily used the search function to find the vases. A couple of subjects tried to answer the task questions without even browsing the vases. A couple of subjects tried to abandon Perseus and asked for permission to use google.com to search the vases directly. Many insisted on typing the questions as-is into the system and were very disappointed when they failed to get a summary answer. Nearly half of the group tried to open too many windows, slowing down the system even more. The P subjects were particularly impatient, tending to stop image loading half way.

7.6.1 TASK

Out of the ten task questions in Appendix C.2, the V group scored 5–8 with a variance of 1.07; whereas the P group scored 4–9 with a variance of 4.8. A summary of the task results was tabulated in Appendix D.2. The averages were shown in FIGURE 7.1. Task

ANOVA was given in TABLE 7.1 with a ttest p-value of 0.007379. The task results show that the V groups (averaged 7.2 correct) performed better than the P group (averaged 5.4 correct).

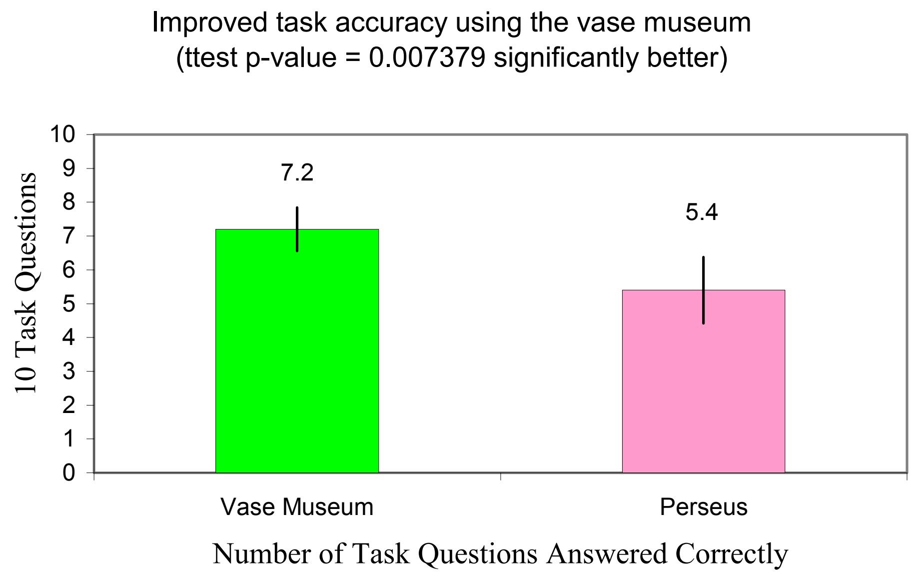

FIGURE 7.1 TASK PERFORMANCE

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
V	10	72	7.2	1.066667
P	10	54	5.4	2.488889

ANOVA

Source of Variation	*SS*	*Df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	16.2	1	16.2	9.1125	0.007379	4.413863
Within Groups	32	18	1.777778			
Total	48.2	19				

TABLE 7.1 TASK ANOVA SINGLE FACTOR

7.6.2 TASK TIME

The users were timed for each of the ten task questions. Time used in the vase museum was spent in selecting the view ports and navigating the scene by the movement keys. Time used in Perseus was spent in clicking the web links of the vases and waiting for the pages to load. Most subjects tried to view more than one vase before selecting an answer. As shown in FIGURE 7.2, the V group was able to complete the task in an average of 13.69 minutes, as compared to the 37.03 minutes for the P group, in nearly 1/3 less time. Time ANOVA for the task was given in TABLE 7.2 with a ttest p-value of 0.000114. The task time was tabulated in Appendix D.2.

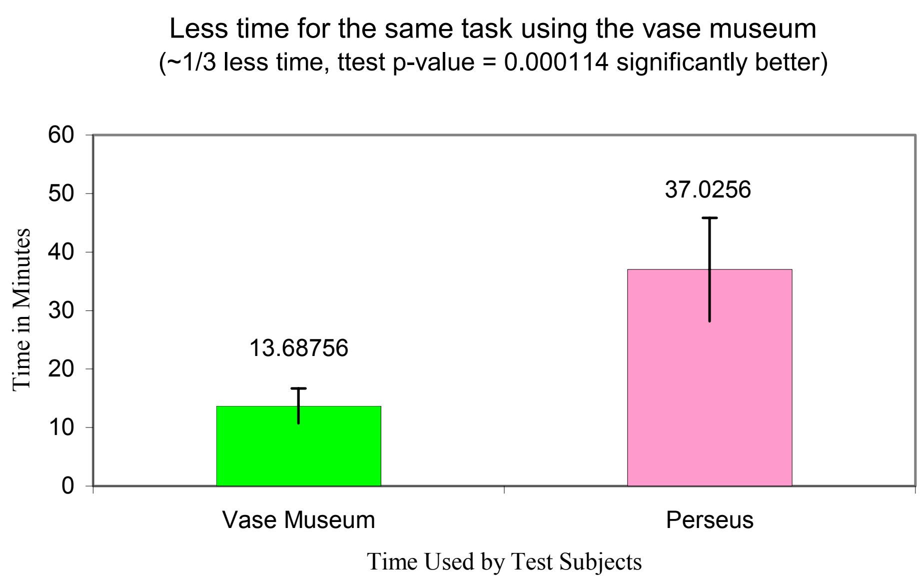

FIGURE 7.2 TASK TIME PERFORMANCE

SUMMARY

Groups	Count	Sum	Average	Variance
V	10	136.8756	13.68756	23.31957
P	10	370.256	37.0256	203.0564

ANOVA

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	2723.321	1	2723.321	24.06015	0.000114	4.413863
Within Groups	2037.384	18	113.188			
Total	4760.705	19				

TABLE 7.2 TASK TIME ANOVA SINGLE FACTOR

7.6.3 TASK SURVEY

Since there were eight task survey questions associated with each task question, the grand averages of each task survey question (S.#.1 .. S.#.8) were taken for the V and P groups, respectively. The task survey can be found in Appendix C.3 and the averages and grand averages for both groups can be found in Appendix D.3. Survey questions S.#.7 and S.#.8 do not apply, and therefore were not administered, to the warm-up question or the first task question. The averages of S.#.1 .. S.#.8 were calculated for each subject (# = 1..10) in the group (V or P), and then the averages of all subjects for the group were calculated to get the grand average for the group. FIGURE 7.3 on the next page indicated little difference in the results of the task survey questions. S.#.3, S.#.6 and S.#7 were practically the same, showing no preference to either system. In S.#.1, S.#.2, S.#.4 and S.#8, subjects preferred

Perseus, but not significantly. S.#.5 was the only one that the V group was significantly better than the P group (ttest p value = 0.036292).

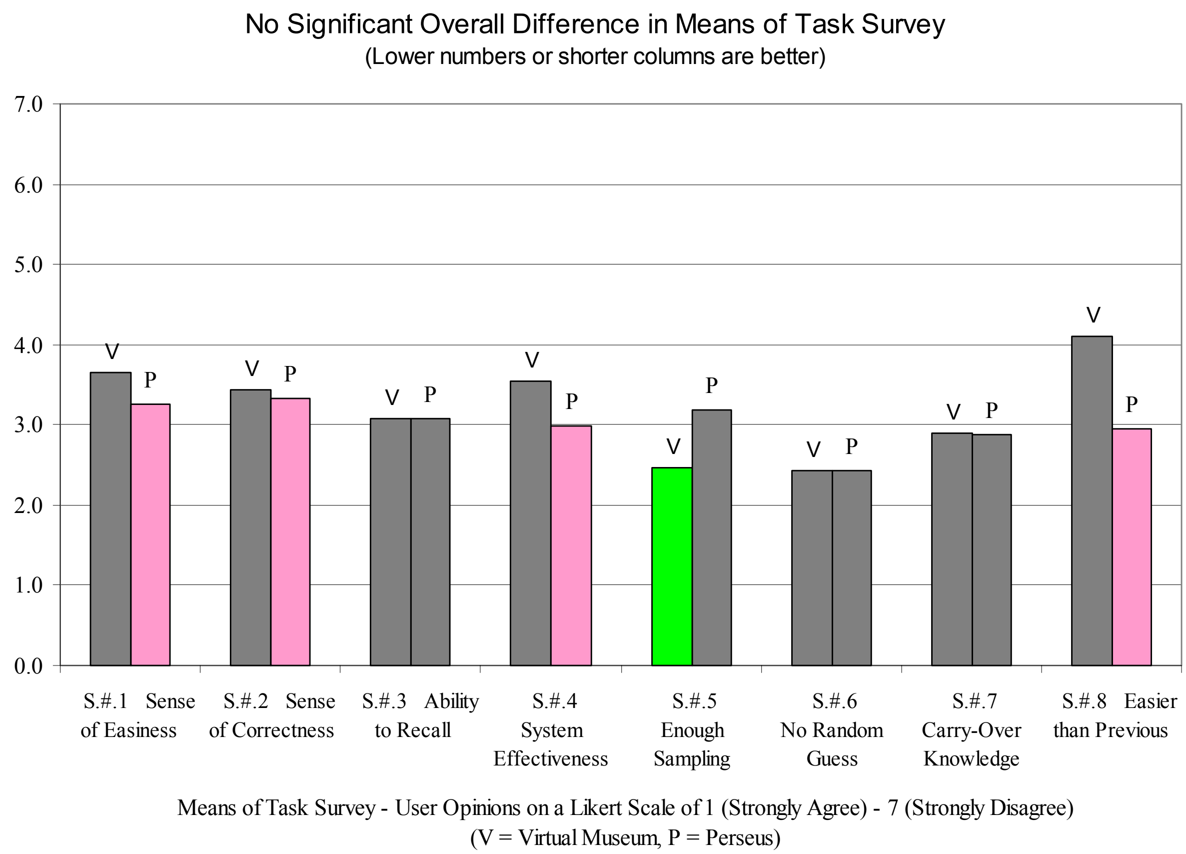

FIGURE 7.3 TASK SURVEY RESULTS

7.6.4 QUESTIONNAIRE SURVEY

A shown in FIGURE 7.4 on the next page, out of the 20 (A.1 … A.20) survey questions in Part A of the questionnaire, there was no overwhelming difference in subject preference to either system. Part A of the questionnaire can be found in Appendix C.4.1 and a summary table of means for Part A of the questionnaire can be found in Appendix D.4.1.

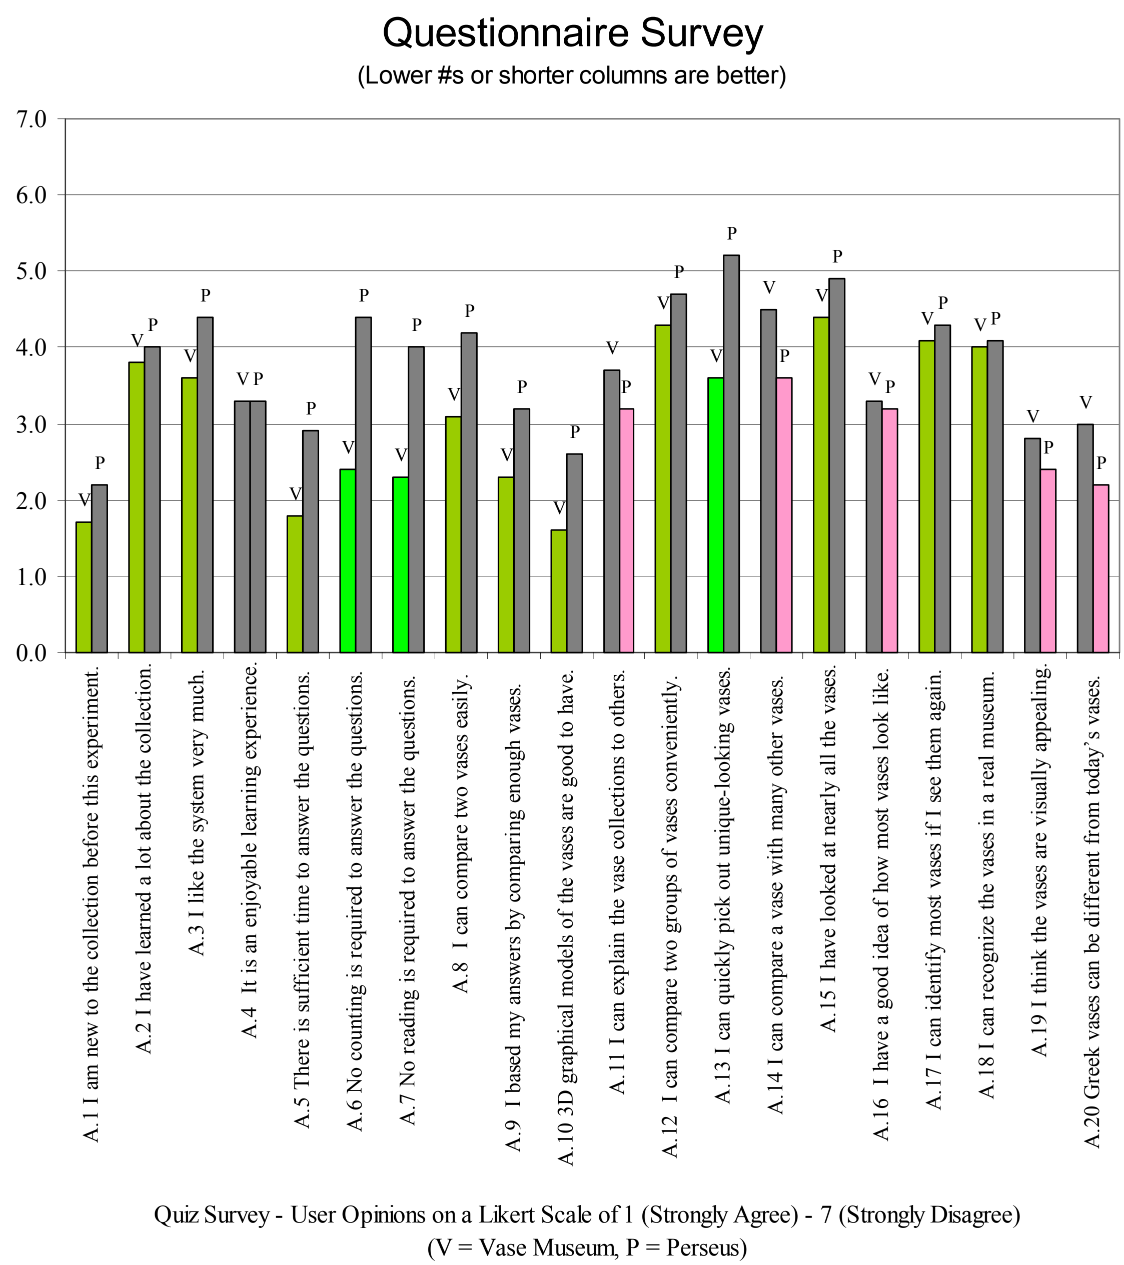

FIGURE 7.4 QUESTIONNAIRE SURVEY RESULTS

In A.16, A.19, and A.20, subjects preferred Perseus, but not significantly. A.4 was a tie. In A.6 (ttest p-value = 0.0125), A.7 (ttest p-vale = 0.0351) and A.13 (ttest p-value = 0.0381), subjects significantly preferred the vase museum. The remaining survey questions showed a slight preference for the vase museum, but not significantly.

7.6.5 QUIZ

As shown in FIGURE 7.5, the V group scored an average of 36.1 (out of 50) in mean time of 13.69 minutes; whereas the P group scored an average of 33.5 in mean time of 37.02 minutes (Section 7.6.2, Appendix C.4.3). Quiz ANOVA was given in TABLE 7.3 with a ttest p-value = 0.3750, which favors V but not significantly.

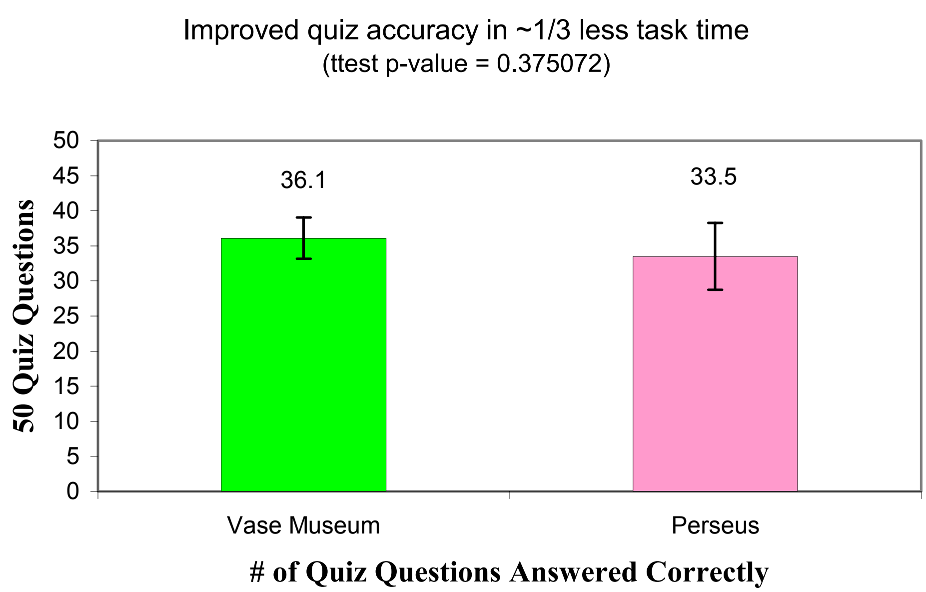

FIGURE 7.5 QUIZ PERFORMANCE

SUMMARY

Groups	Count	Sum	Average	Variance
V	10	361	36.1	22.54444
P	10	335	33.5	59.16667

ANOVA

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	33.8	1	33.8	0.827305	0.375072	4.413863
Within Groups	735.4	18	40.85556			
Total	769.2	19				

TABLE 7.3 QUIZ ANOVA SINGLE FACTOR

7.7 DISCUSSIONS

Our testing results were encouraging, showing significant improvement in task accuracy (p=0.007) and speed (p=0.0001) in favor of the V group. All subjects were able to complete the task. Out of the ten task questions, V subjects scored 5–8 with a variance of 1.1; whereas the P subjects scored 4–9 with a variance of 4.8. The V group was able to complete the task in an average of 13.69 minutes, as compared to the 37.03 minutes for the P group, nearly 1/3 less time in favor of the V subjects. The V subjects were more able to answer the questions based on looking at many vases (ttest p-value = 0.036). They feel that they were more able to pick out a unique looking vase from the collection (ttest p-value = 0.038). They also agreed more to the statements that no reading (ttest p-value = 0.0125) or counting (ttest p-value = 0.0351) was necessary to perform the task. We also see significant improvement in quiz performance, where the V group learned just about the same amount of information in nearly 1/3 the amount of time. We also found that subjects preferred more contextual visualization, greater navigation ease, better search ability and faster system speed. Neither group showed a noticeable preference to the system in use.

Both groups complained that the systems were too slow, but in different ways. The P group was impatient because they needed to wait for the HTML pages to be loaded one vase at a time. In trying to answer a collective question, this can be frustrating since the system cannot load all the vases. Therefore, the P group was compelled to open a handful of windows, view the pictures, close them, and open some more. Given the number of questions in the limited timeframe, most users in the P group complained that they did not have sufficient time. They believed they would do better if the system was faster or if they could see more pictures. The V group, on the other hand, was frustrated because 'walking' was an ineffective method of access. The system was designed such that the users could navigate in the room and approach (by walking to) vases of interest. Although this held true in the testing, the users seemed to demonstrate a different perception of presence. It appeared that if they were able to see a vase in the display (on the screen), they believe they were already there. Therefore, they were impatient in having to hit the cursor keys so many times just to 'get there.' Many users would first click on a vase or drag a region around it in an attempt to 'navigate' to it. For example, a few users would click on the vase in the scatter plot first in the hope that they would be immersed in the scene as a result.

The V subjects were new to desktopVR and were more tuned in for warm-up training. The P subjects were seasoned web users who would gloss over initial directions. For example, a couple of P subjects preferred not to left click on the links; instead, they would always right click and select from the drop down menu. P subjects who were accustomed to searching would go straight to the search field even though the correct link was already displayed right in front of them. Since the task questions were characteristic of most vases in the collection, the P subjects would have been able to answer correctly if they have selected a few vases. A few questions could have been answered right off the vase list page. Unfortunately, searching in the entire library returned a list that included all collections, and not just the London vases. The skill set of the P subjects could be an over-kill to the task and served to their detriment.

We also observed that the V subjects tend to arrive at their answers by inclusion. They would navigate and observe enough vases in the background to justify their answers. For example, one of the task questions asked them for a color that they would not find in the collection. The P subjects would determine the colors of the vases by the ones they have already seen in the 3D scenes. The P subjects tend to arrive at their answers by exclusion. Since they could not have seen enough vases in one glance, they would randomly select vase HTML pages and cross off answers that would not apply. They would continue to do so until they have only one choice left. The P subjects would take a long time using this method of exclusion if they failed to exclude enough choices. This occurred frequently because they have a tendency to sample vases that were grouped together in the list. Since the vases were sorted by catalog numbers, the order of the list bears no significance in the context of the collection.

CHAPTER 8

THESIS CONCLUSION

The issues we are addressing in this thesis are characteristics of the focus versus context problem as applied to the digital library. We have recognized the crippling effect in focus versus context. By removing the crippling effects, the solution we have proposed is the invention of new interaction techniques for 3D focus+context interactive visualization of a large data set. The work in this thesis requires a merge of human-computer interaction, information visualization and virtual reality. Within the domain of human-computer interaction, the prototypes concentrate specifically in areas pertaining to next-generation NON-WIMP interaction techniques in desktopVR, expandable to immersiveVR. Within the domain of information visualization, the prototypes present selected interactive visualization techniques applied to information extracted from Perseus. We would like to stress the point that some of the solution concepts cannot be cost-effectively implemented given accessible virtual reality speed. However, we have built three major prototypes and many experimental models to demonstrate the technical and functional feasibility. We have presented a prototype environment that can be run on a laptop with or without immersive equipment such as a head-mounted display. The prototypes we have built showed that desktopVR and even immersiveVR are robust enough, and scaled-down versions and 3D models can be used in existing digital library systems.

Because of our work in this thesis, we have reached the following conclusions:

- The focus and context problem in the digital library can be solved by reducing it to the problem of crippling effects.
  - Using 2D text to describe a 3D world is a form of context crippling.
  - Using 2D images to represent 3D objects is a form of data crippling.
  - 2D information crippling is the sum of 2D context and data crippling.
  - In virtual reality, 2D information crippling becomes 3D virtual crippling.
- The crippling effects can be overcome in virtual 3D.
  - Context crippling can be solved by 3D virtual reconstruction of the 3D world.
  - Data crippling can be solved by 3D virtual reconstruction of the 3D objects.
  - 2D information crippling can be solved by 3D interactive landscape populated with 3D data objects in contextual visualization.
  - 3D virtual crippling can be solved by solving 2D information crippling.
- Information in focus+context can be more effectively and efficiently delivered in virtual reality
  - Users can learn more accurately and in less time about visual information by ‘interacting’ with data objects in a 3D world than by ‘reading’ 2D text images on a 2D screen.

The 3D information space proposed in this thesis presents a data world where information is delivered in its context. The techniques introduced restore the 3D nature of data artifacts and focus on new ways of visualizing and communicating contextual information in VR. To prove feasibility and usability, we have built three major prototypes to demonstrate the new set of interaction techniques in desktopVR, expandable to immersive VR. We have created 3D models of data objects from actual information found in Perseus. The first prototype is a London walkthrough where we present a map overlay of plot, building and street maps. The second prototype is a scatter plot of a Greek coin catalog populated by data coins. The third prototype is a Greek vase museum with 3D vases situated in a scatter plot. In all the prototypes, the users can view realistic data and navigate interactively in an information landscape. The focus of interest is determined by proximity to specific data while exploring the entire context of the collection. In a usability test, we compared one of the prototypes we have created, the vase museum, to the equivalent site in Perseus. We have two groups of 10 users each. One group learned the vase collection by browsing in Perseus. The other group learns the same vase collection by using our vase museum prototype. We have designed a set of task questions and a questionnaire that measured the speed and depth of learning. Our usability test showed that users who used our prototype learned more about the same collection of digital documents in less time.

APPENDIX A

PROTOTYPING ENVIRONMENT

The three prototypes presented in this thesis are populated with a significant number of 3D models. There are 242 London buildings in the London walkthrough (Chapter 4), 132 Greek coins (Chapter 5) and 157 Greek vases (Chapter 6). In addition to these primary 3D models of ancient artifacts, there are also many secondary or supporting 3D models in the virtual scenes, such as rooms, floating HTML pages, columns, 3D words, and so on. The 3D models are created from photographic images or text descriptions found in Perseus. Many parts of the virtual scenes require complex treatments such as bending the streets and the vase handles. The following subsections describe the general prototype environment shared by the three prototypes.

3D COMPONENT PROTOTYPING. To increase productivity and reduce cost, we employ techniques in 3D component prototyping. The virtual scenes are created by an industrial-strength 3D graphics package, exported to executable codes and rendered by third-party loaders and plug-ins. 3D component prototyping allows us to experiment our concepts rapidly in the implementation process, freeing us to concentrate more in the designing aspects rather than in the programming aspects. The selection of the 3D components used in the prototyping process is based on empirical reasons, such as the robustness of the system, the number of texture maps allowed and loading speed of a large virtual scene.

The virtual scenes presented in this thesis are modeled in **Discrete© 3ds max 4.3** (MAX). [57] Objects in the scenes can be imported and exported to other graphics or programming platforms. Imports are limited to pre-made 3DS objects or other VRML (Virtual Reality Markup Language) scenes. [51, 201] After modeling completion, the entire scene is exported to **VRML97** file format. Only JPEG file format is used for the texture maps because it is recognized by all VRML97-compliant browsers. All JPEG files are preprocessed in **Microsoft™ PhotoDraw© V2**. JPEG files are created from photographic images downloaded from Perseus. The entire prototype system can be operated on a PC laptop, eliminating the need for a lab environment. Video productions of the VRML scenes are captured in a time-limited, share-ware version of **Advance-Visual Interface** (AVI) using **Camtasia Studio**© offered by **TechSmith™**.

PROGRAMMING PLATFORM. Interfaces for the immersive equipment are written in either C or Java [101] and are called in the VRML scenes as functions. [139] The prototypes presented in this thesis blur the line between desktopVR and immersiveVR even more by presenting high-fidelity mock-ups supportable in both kinds of virtual environments (VEs). Currently popular programming platforms for VEs are C/C++, Java/Java3D, and VRML on PCs, Silicon Graphics, or other proprietary systems. Regardless of the language or systems chosen, coding VEs from scratch can be a formidable task. To cut development time, it is highly desirable to have access to modeling sources with pre-built-geometries that can be imported into VEs with the proper loaders. **Viewpoint Datalabs International** has a database of over 10,000 3D models. **Superscape's DO 3D** and **VRT** systems offer a significant warehouse of pre-built objects. Coordinates of 3D models can be also generated by a **CAD/CAM** and converted by filters to polygonal forms to be imported into a VE. A myriad of modeling software can also be used to construct objects in a VE. These systems can cost anywhere from hundreds to thousands of dollars. Almost all require a solid understanding of computer graphics and a steep learning curve. [67] Many immersive virtual environments (IVEs) are constructed from simple geometric primitives, sacrificing realistic graphics for faster rendering. Graphic realism may also be discounted due to the restricted resolution of the head-mounted displays (HMDs).

FILE TRANSFERS. Import and export functions allow 3D models or even the entire scene to be exchanged among different development platforms such as 3ds max™, **Maya™**, **Mechanical Desktop**, **Pro/ENGINEER**, **SOFTIMAGE®**, **CATIA®**, etc. Some popular geometry file formats are: ASCII (ASE), **Alias/Wavefront** (OBJ), 3D **Studio Mesh** (3DS), **3D Studio Project** (PRJ), **3D Studio Shape** (SHP), **Adobe Illustrator 88** (AI), **AutoCAD** (DWG), **AutoCAD** Interchange (DXF), **Initial Graphics Exchange Specification**[1] (IGE, IGS, IGES), **Stereolithography** (STL), and VRML (WRL, WRZ) files. Unfortunately, not all features provided in the different programs are freely exchangeable and sometimes special loaders or scripting may be required.

[1] IGES is an ANSI standard for the exchanges of NURB data.

RENDERING. The prototypes in this thesis are rendered by the **Cortona VRML** Browser by **Parallel Graphics**.[2] After a VE is modeled, it must be rendered. 3D rendering generates perspective images of the VE as will be seen by the avatar. Before an image is displayed, it is formed inside the computer's memory by the renderer program. While this process is underway, the previous image is being displayed. This overlapping of computer tasks is another method of keeping the VR system running as fast as possible. The renderer uses the perspective view of an object to determine the pixels to update. It then uses the color of the object to flood the relevant pixels. The image can be made to look life-like by coloring the object with color shades that give the impression that it is lit by some light source. The renderer can also simulate lighting by working on the surface shading of the object. To do this, the renderer is given the position and intensity of a virtual light source, and using some simple laws of illumination, shades the object with acceptable color intensities. In our prototypes, we simulate surface lighting by touching up the texture-maps with contrasts and image sharpening tools. The rendering is technically separate from 3D modeling. However, high-speed rendering is vital to a successful VR system. The illusion of immersion is compromised if there is any delay in producing the images seen by the users. Consequently, the majority of VR systems are built around VEs modeled from polygons or triangles. Between the two, triangles are preferred as they are consistent, i.e. there are 3 sides and they are always flat. A polygon built from 4 or more sides can be twisted, which can introduce errors in rendering. Since we are using 3ds max 4.3 for 3D modeling and Cortona

[2] A version of Cortona VRML Client can be downloaded from:
http://www.parallelgraphics.com/products/cortona/download/iexplore/

VRML Browser for 3D rendering, it is very important in the prototyping process to use features that are supported by both packages. For example, the simple 'transparency effect' (achieved by setting the transparency factor) is exportable from 3ds max 4.3 to Cortona VRML, but the 'glass effect' (achieved by creating a 'material' in a material library) may not be exported properly --- even though the outputs may appear indistinguishable to the viewers (both are 'see-through').

PRIMITIVES. Primitives are 3D graphical objects populating a VE. Canned primitives can be created by programming function calls, procedure calls or instances of classes. They can also be created using a 3D graphic package (such 3ds max 4.3). Depending on the platform, the basic primitives in 3D modeling may include varies boxes (cubes and rectangles), spheres, cylinders, cones, pyramids, prisms, lines and planes. Some systems may offer more advanced primitives such as toruses, tubes, capsules, arrows, hedras, spindles, L-gons, C-gons, Gen-gons, particles, teapots, splines, text, or even bones and skeletons. Primitives can have attributes that define the vertices, color, location, rotation and scale factors. A group of simple primitives can be assembled to form a more complex object. For example, a table can be created with a large, flat rectangle on the top with four long, skinny rectangles as legs. Materials such as wood grains can be applied to the desk to give it a more realistic look. Materials are texture maps (image files) that are pasted onto the facets of the primitives. **UV mapping** (coordinate mapping) can be used to control the texture-maps of specific regions. Techniques in **tiling** can be used to repeat a texture-map pattern, such as a parquet floor in a room.

The buildings in the London walkthrough are modeled from the box primitive. The Greek coins in the coin catalog are modeled from two oil tanks (bulging disks). An alternative solution to the coins can be one oil tank with two texture maps for the front and back of the coins. The semi-transparent HTML pages in the vase museum are very thin, texture-mapped boxes (so that they appear to be a 'page'). The vase captions and HTML pages are both **billboards** facing the direction of the users so they can be seen from all angles. The Greek vases are very complex meshes created from lathes of 2D line drawings. Simple arch-like handles are created from bending toruses or cylinders. Elaborate handles must be made from complex meshes.

Triangulation of the meshes seems to be most stable in exports. There can be a creative limitation to building complex objects by assembly. The more advanced objects in most decent 3D worlds are comprised of meshes. Meshes are highly detailed geometric grids or triangulation of 3D objects defined by the object's vertices. Each facet of the mesh surface, the polygons, the vertices and the edges can be manipulated independently. The exact coordinates of mesh vertices can be exported from a Computer-Aided Design (CAD) system and imported to a 3D scene. Simple primitives such as a cube can be converted to a mesh with 6 facets (top, bottom, left, right, front, back) or more (a $n \times m$ matrix). It is not uncommon to have a mesh model consisting of thousands of surface polygons.

MODIFIERS. Modifiers are **deformers** that allow the users to sculpt and fine-tune the objects into any desired odd-shaped forms. Modifiers can be **parametric** deforming values such as the locations, angles and amount to bend, twist, skew, taper, wave, warp, squeeze, stretch, exclude and so on. For example, the simplest way to model an obelisk is to create a rectangle, texture map it with an image of the structure, and taper it at the top. Modifiers can also be **free-form** deformers, which allow advanced editing on complex meshes such as facial expressions. Orientations, translations, colors and materials are attributes of the primitives and not of the modifiers. Modifiers respect the hierarchy established by groupings or scene graphs. The position in the hierarchy where modifiers are applied in a scene graph is very important because all the children of the deformed node will be affected. Modifiers are definitely an area where a using 3D graphical modeling package surpasses the conventional programming approach.

LIGHTS. Without lights the VE will be pitch black. For simple VR applications, scenes illuminated with default lights suffice for most general purposes. Default lights illuminate the entire scene uniformly without a specific light source. There are times when default lights fail to provide the realism needed. Lights can be used to create the illuminations, reflections, refractions and shadow effects generated by sources such as lamps, sun, etc. The types of light sources commonly available are spotlights, directional lights, targeting lights and omni-lights. In an IVE, given the resolution of currently affordable HMD technology, advanced lightings are over-kills. However, because VR worlds are black by default, omni light sources other than the default lights may be needed to highlight areas of interest. Users may not be tempted to navigate toward a dark area where there is no scene. It is possible to have different colors of lights, but in a VE, where the colors of the modeled scenes should dominate, white lights should prevail.

CAMERAS. Cameras define the viewpoints, which is what a viewer sees in a rendered scene. Without preset camera views, the entire scene is shown as a whole, which may be too big to be visually meaningful. Cameras are analogous to the photographic lens used in the real world. Cameras focus the lights reflected by a scene onto a focal plane. The distance between the camera lens and the focal plane is the focal length of the lens. Focal length determines the field of view (FOV), which is the area of the scene rendered. The focal lengths are measured in millimeters (mm) and the FOV is measured in degrees of the horizon. The standard focal length in common photography is 50 mm with a 46 degree FOV. The perspective offered by 50 mm lenses appears to be most normal. Short or wide-angled lens have focal lengths shorter than 50 mm with wider FOVs, which can show more of the scene in greater perspective. Wide-angled lens are to be avoided in 3D modeling of an IVE due to the problem with real-time rendering speed. The more of the scene to be rendered, the more time it takes and the more of a delay the users may feel, which may reduce the overall sensation of immersive realism. Lenses with a focal length longer than 50 mm and narrower FOVs are called long or telephoto lenses. They show less of the scenes, but provided more detail, as rendered objects appear to be flattened and parallel to the viewer.

STEREOSCOPIC VISION. Ideal images in an IVE when viewed through an HMD should be made stereoscopic, or one image per eye, in order to simulate the depth cues provided by the two human eyes. To achieve stereoscopic vision, two scenes of the IVE are generated, one for each eye, offset by the ½ of the distance between the pupils from the gaze, which may differ from person to person. An image that had been treated stereoscopically appeared out of focus (blurry) on a standard computer monitor, but appeared in focus (clear) when viewed through a HMD. Unfortunately, stereoscopic vision doubled the amount of graphics the computer had to load. In preference for faster rendering speed, the scenes in the thesis prototypes are monoscopic, which meant there is no depth information. Depth is perceived through perspective modeling of the VE as well as the user visual and motion adjustments. The London walkthrough supports a HMD with head tracking. The coin catalog is too slow for HMD to work in real time even though it is implemented. Instead, animations are used to rotate the coins so users can get an all-around view. The vase museum is too complex and slow to be an IVE. In general, IVE is still a work in progress for all the prototypes.

SCENE GRAPHS. Scene graphs are the hierarchy of the objects in VE. They are generally represented in the design phase as a tree with the world as the root node and the objects in that world as paths of children nodes. Conceptually, they are stacks of grouped objects. Without a well-structured scene graph, the VE would be populated with a myriad of independent primitives without a sense of cohesion. For example, to model the hand, the parent node is the hand with five children nodes, which are the thumb and the index, middle, ring fingers and the pinky. To move the entire hand with fingers and all, the translation matrix can be applied to the parent node in the scene graph. Without this hierarchical organization of objects, it may be a challenge to move each finger individually while maintaining the overall modeling structure of a hand.

In the London walkthrough, the root node is the intersection. The intersection can be broken down into ground, lights, background, streets, etc. Each street is broken down into the buildings. Each building is broken down into six facets, each with a different texture map. In the coin catalog and the vase museum, the root node is the room (with walls, ceilings and floors) and the internal nodes are the lights, cameras, artifacts (buildings, coins, vases), etc. The coins are broken down into an animation node with two oil tanks. The vases are broken down into the actual vase meshes, plus all supporting objects such as HTML pages, captions, billboards, lights, cameras, etc.

X-REFERENCE. An X-reference, or simply X-Ref, is an externally referenced file or modeled object in Discrete© 3ds max 4.3. A virtual scene can contain X-References to other scenes or other objects in an external file. X-references allow multi-tasking of a large scene without interference among the developers. Several modelers or animators can work on their parts of the final scene in their own separate files.

In the vase museum, the vases are created individually in a separate file. All vases are hidden at the origin (0,0,0). Only one current working vase is visible at a time. Six camera views showed the left, right, front, back, top and bottom views of the vase and handle meshes. When the file is exported, only the current vase is exported because the other vases are hidden from view. This allows individual vases to be quickly exported using only one file. The individual vases can be launched from the corresponding vase HTML page in the virtual scene. In each of the renditions of the vase museum, the individual vases are imported as X-References in the scene. This allows many versions of the vase museum to be quickly prototyped without affecting the original vase models. A change in the original vase file will automatically be propagated in all the scenes that call the vases by X-References. However, each rendition of the vase museum must be re-exported to reflect the changes in X-References.

SYSTEM REQUIREMENT The computer requirement is Pentium 90 MHz or better, running Microsoft® Windows 95/98/ME/2000/XP or Windows NT ® 4.0. The browsers required are Internet Explorer 4.0 or a later version, Netscape®, Navigator 4.0 or a later version, or Mozilla 1.0 or a later version. A minimum of 16 MB of random access memory (RAM) is required. The display must be SVGA with 800x600 high color mode and higher. The actual prototypes are run on a Toshiba® Satellite 5005-S507 Intel® Pentium® III Laptop with 1.10GHz, NVIDIA® GeForce4 graphic card, and 32 MB of VRAM.

PERFORMANCE. All virtual scenes created in the prototypes are optimized for VRML97 exports.[3] [196, 199] Animations other than position, rotation and scaling are avoided. Lights supported are direct, omni and spotlights. Cameras can be free or targeted. An optimal scene should contain a maximum of 5,000–10,000 polygons with minimal use of texture maps. [77, 79, 165] To reduce file size, primitives, rather than complex meshes, are used whenever possible. The performance of VRML files can be improved if the users are restricted to a limited window on the screen to reduce the number of pixels that must be rendered on every frame. The embed command (for example, WIDTH=300, HEIGHT=400) can be used to improved performance by setting a smaller window. However, the HMDs

[3] The complete VRML97 specifications can be found at http://www.vrml.org/Specifications/VRML97

can be run only in the full-screen mode. Performance can also be improved by using instances of same geometry, as opposed to new copies. The rendering frame rate specifies the processor usage in the range from 0 (idle and minimum) to 100 (maximum). All the prototypes are run at the maximum frame rate. Navigation may appear to be 'chunky' because frames are being dropped as the browser strains in real-time rendering. Navigation can be ineffective in a large virtual space, where the users can see the target, but can take a long time walking up to it (by hitting the cursor key down continuously).

APPENDIX B

PROTOTYPE SOURCE DATA

B.1 ORIGINAL TALLIS MAPS

FIGURE B.1 is a map of the long Leadenhall Street that starts at the right side of the top two rows. It goes from left to right and continues from right to left on the bottom two rows. A corresponding street map is shown to the right of the building rows. An enlarged image of a building on the street is shown to the left of the building rows.

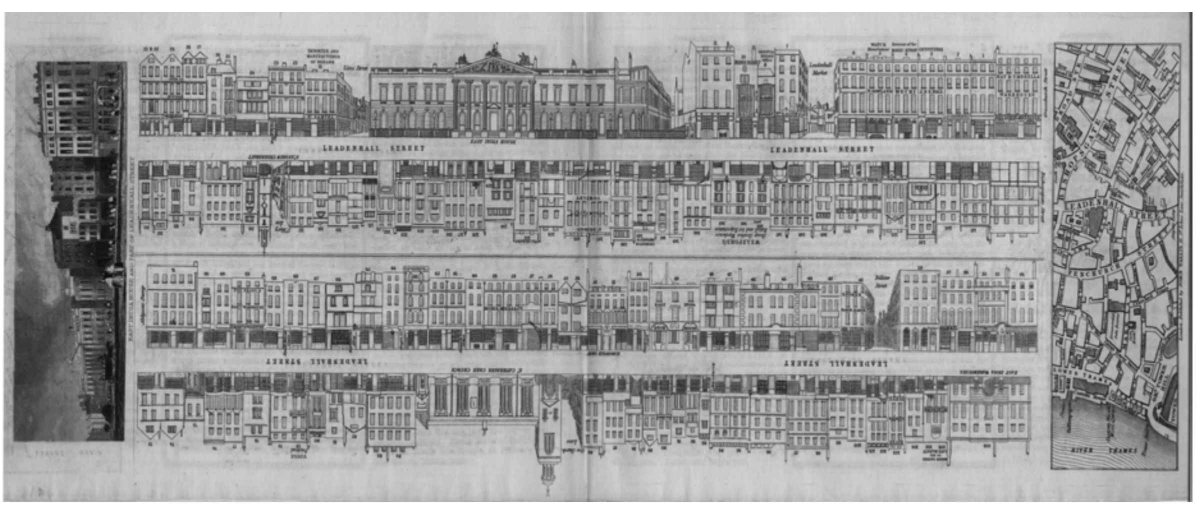

FIGURE B.1 TALLIS MAP OF LEADENHALL (LONDON 1838-40)

FIGURE B.2 is a map of Cornhill (Note: Just 'Cornhill'; No 'Street' or 'Ave' associated with it). This time, the intersection starts at the right side of the bottom two rows. It goes from right to left and continues from left to right on the top two rows. FIGURE B.3 shows the map that completes the intersection.

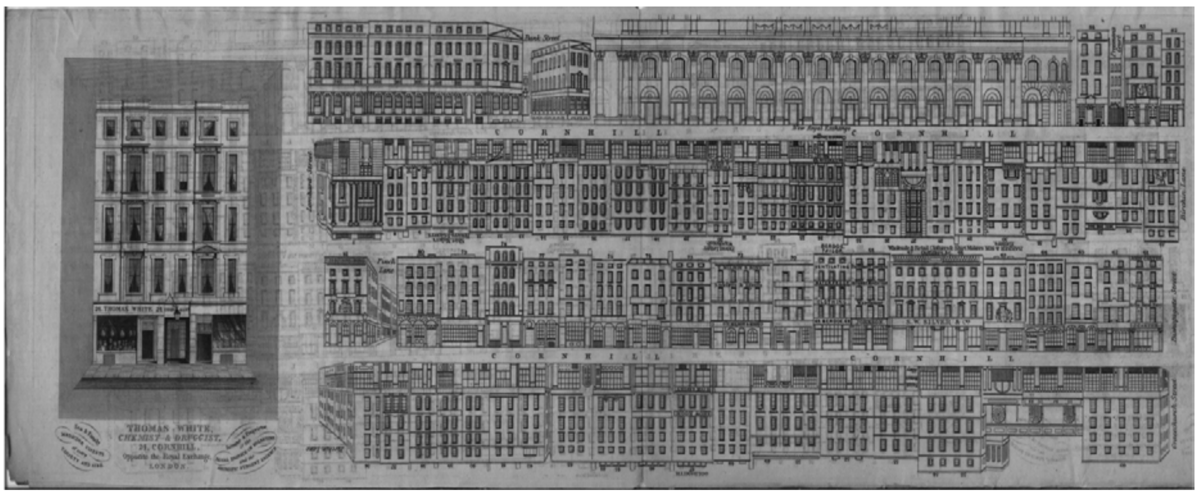

FIGURE B.2 TALLIS MAP OF CORNHILL (LONDON 1838-40)

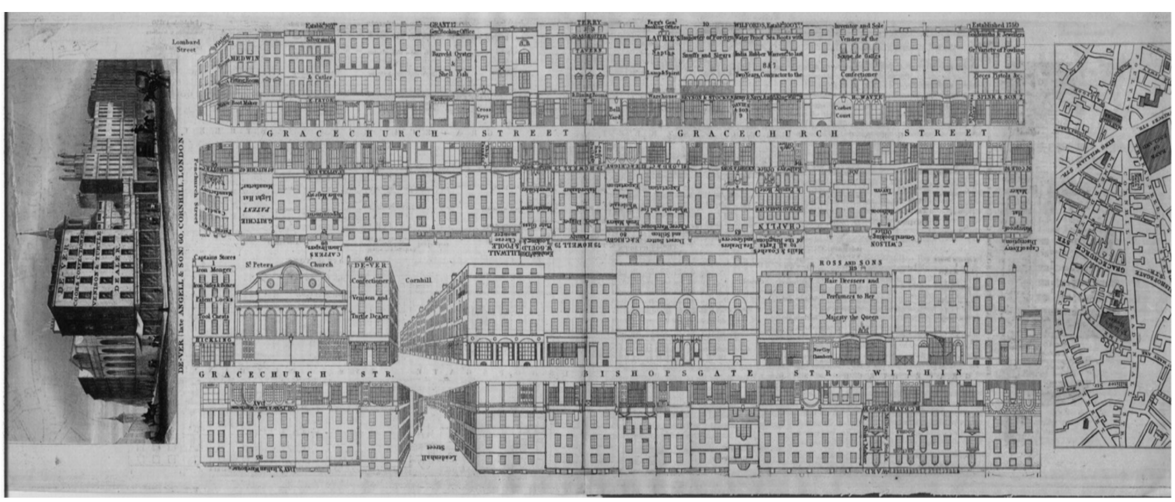

FIGURE B.3 TALLIS MAP OF GRACECHURCH AND BISHOPSGATE (LONDON 1838-40)

FIGURE B.3 contains two very short streets: Gracechurch Street and Bishopsgate Street. On this map, the streets start from Gracechurch Street at the lower left drawing where an intersection with Cornhill and Leadenhall Street are shown. The streets span out from that intersection point. To the right is the very short Cornhill. To the left is Gracechurch Street, which continues from left to right on the top. Just like FIGURE B.1, the street map is shown to the right and an enlarged image of a building on the street is shown to the left.

FIGURE B.4 shows a partial plot map of Bishopsgate Street and FIGURE B.5 shows a partial plot map of Cornhill. These two maps provide for the missing depth perception of the building lots. They also show more buildings than what are shown on the previous three maps.

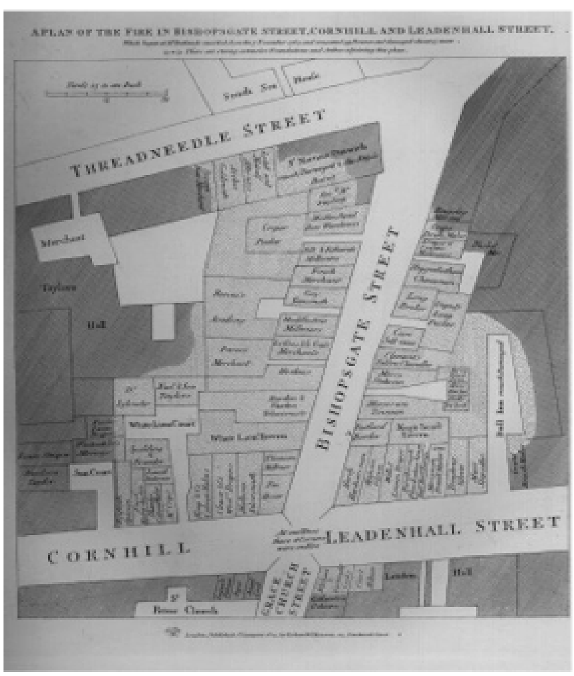

FIGURE B.4 PLOT MAP OF BISHOPSGATE

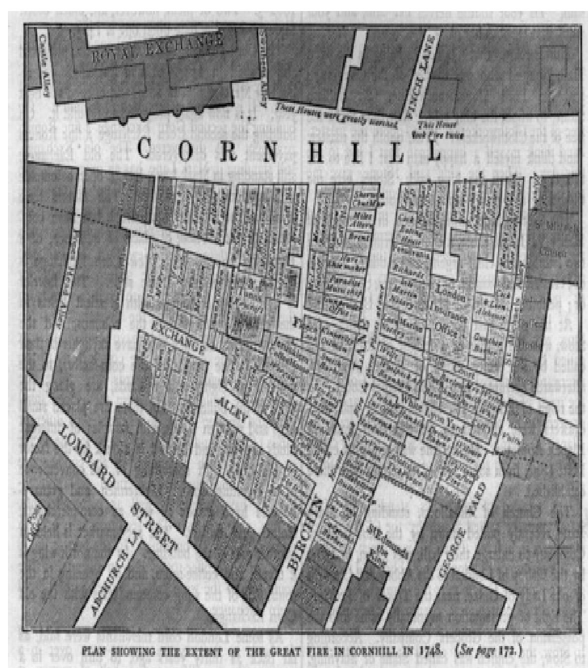

FIGURE B.5 PLOT MAP OF CORNHILL

FIGURE B.6 (extracted from the right part of FIGURE B.1) shows the street map of the BCGL intersection. This map is used as the 'ground' JPEG (JPG) texture map and is only visible from the bird's eye view.

FIGURE B.6 TALLIS MAP OF THE BCGL INTERSECTION[1] (LONDON 1838-40)

Collectively, Figure B.1 to B.6 form the group of relevant Tallis maps used to reconstruct the 3D interactive London walkthrough described in Chapter 4. There were 242 buildings modeled for the intersection.

B.2 TEXT DESCRIBING THE LONDON FIRE OF 1765

A London fire occurred at the BCGL (see previous section B.1) intersection in 1765.

[1] The red asterisk is intended to show the center of the BCGL intersection and is not part of the original map.

The following text passage available in Perseus[2] is used as the story setting for the 3D text menu described in Section 3.4.3 and Section 4.4. It is the complete text version of the HTML page shown in FIGURE 5.4 in Section 4.4.1. Note that B.4 is shown as the supporting map to the text passage.

<Beginning of Text>

Plan of the Fire in Bishopsgate Street, Cornhill, and Leadenhall Street: November 7th, 1765.

Plan of the Fire in Bishopsgate Street, Cornhill, and Leadenhall Street: November 7th, 1765.

A Plan of the Fire in Bishopsgate Street, Cornhill and Leadenhall Street.

Two of the most fatal conflagrations which have occurred in London since the entire destruction of the City, have broken out within a very short distance of each other; the devastation of the second continuing, as it were, the ravages of the first. Of the former of these some account is given in the note beneath,[a] and the latter is the subject of the present pages; while the annexed Plate, originally published within a few days after the melancholy event,[b] exhibits the spot where it began and the extent of its destruction.

This terrific Fire first broke out soon after three o'clock in the morning of Thursday, November 7th, 1765, at the house of William Rutland, Peruke-maker,

[2] Robert Wilkinson, Londina Illustrata: Graphic and Historical Memorials of Monasteries, Churches, Chapels, Schools, Charitable Foundations, Palaces, Halls, Courts, Processions, Places of Early Amusement, and Modern Present Theatres, in the Cities and Suburbs of London and Westminster: Volume 1
At the time of this thesis, this article in Perseus can be found at:
http://nils.lib.tufts.edu/cgi-bin/ptext?doc=Perseus%3Atext%3A4000.01.0144%3Aid%3Dc.6

marked A on the Plan, situate the second building on the east side of Bishopsgate Street Within, opposite the White Lion Tavern. At this time the wind was high in the south-west, and the flames soon spread across the way, and set fire to the residence of Mrs. Thomson, a milliner; and as it was some time before assistance could be procured, they extended to the corner house kept by Burkuit Fenn, a hosier, whence the other three angles of Cornhill, Grace Church Street, and Leadenhall Street, soon caught the blaze, and were all on fire at the same time. The New River turncock had been summoned immediately on the first alarm, but after having turned off the water in Leadenhall Street, in order to make it flow higher in Cornhill, the two streets being then supplied by one main,--he found that the pipes in Cornhill had been already cut open by the crowd, and that none would rise at the plug where it would have been most serviceable.[c] There was also some delay in the arrival of the office fire engines; the want of which was at first supplied by a large one belonging to Mr. Ephraim Brookes, of Long-Acre, which had been sent to Grocers' Hall in the Poultry, until he completed two for the Company. This was one of the earliest at the spot, and was ordered thither by the Lord Mayor, who sent his own horses to fetch it, immediately on being informed of the fire. It is represented in the papers of the time as having been greatly instrumental in checking the progress of the flames: which is corroborated by a subsequent public advertisement of thanks from the inhabitants of the neighbourhood.[d]

The fire, however, continued still to spread in the most alarming manner. It soon reached the Church of St. Martin Outwich, at the corner of Threadneedle Street, the interior, and much of the walls of which it destroyed; and wholly consumed the steeple, whence the great bell fell down with a prodigious noise. The south and east sides of Merchant-Taylors' Hall also caught fire, and the whole building was saved with great difficulty; but about seven in the morning, the wind shifted to the west, and drove the flames back from the part where they were then raging, otherwise the whole of Great

St. Helen's must have been destroyed, and their course was changed to the north side of Leadenhall Street, in which upwards of twenty houses were consumed. At nine o'clock several parties of guards arrived from the Tower, and soon after Sir William Stephenson, the Lord Mayor, who gave orders for lodging such goods as could be saved in the Royal Exchange, and was otherwise particularly active in affording assistance. From the advertisements of the several sufferers and others, which were published immediately after the fire, it appears that much energetic aid was given to them all, and that many of their goods were saved; but the loss was still immense, and, from the rapidity of the flames, numbers escaped with little more than their lives. The accidents, both at the fire itself and subsequently in the ruins, were very melancholy and numerous, but several of those at first related in the newspapers of the time.

<End of Text>

B.3 DEWING AND BCMA COINS

There are 132 coins modeled in the 3D interactive Greek coin catalog described in Chapter 5. These coins are selected from two collections in Perseus: 62 Dewing ((Dewing Numismatic Foundation) coins and 70 BCMA (Bowdoin College Museum of Art) coins. The coins are identified by the collection and the catalog # (e.g. Dewing 351). TABLE B.1 is a complete catalog list of the 62 Dewing coins modeled in the coin catalog.

#	Coin Catalog #	#	Coin Catalog #	#	Coin Catalog #
1.	Dewing 351	3.	Dewing 499	5.	Dewing 678
2.	Dewing 451	4.	Dewing 611	6.	Dewing 686

#	Coin Catalog #	#	Coin Catalog #	#	Coin Catalog #
7.	Dewing 1007	26.	Dewing 1708	45.	Dewing 2290
8.	Dewing 1023	27.	Dewing 1711	46.	Dewing 2310
9.	Dewing 1289	28.	Dewing 1715	47.	Dewing 2312
10.	Dewing 1309	29.	Dewing 1728	48.	Dewing 2314
11.	Dewing 1496	30.	Dewing 1951	49.	Dewing 2318
12.	Dewing 1501	31.	Dewing 1952	50.	Dewing 2345
13.	Dewing 1533	32.	Dewing 1953	51.	Dewing 2346
14.	Dewing 1543	33.	Dewing 1955	52.	Dewing 2347
15.	Dewing 1545	34.	Dewing 1959	53.	Dewing 2349
16.	Dewing 1546	35.	Dewing 1967	54.	Dewing 2350
17.	Dewing 1561	36.	Dewing 1969	55.	Dewing 2353
18.	Dewing 1567	37.	Dewing 1971	56.	Dewing 2355
19.	Dewing 1571	38.	Dewing 2129	57.	Dewing 2360
20.	Dewing 1573	39.	Dewing 2162	58.	Dewing 2422
21.	Dewing 1654	40.	Dewing 2189	59.	Dewing 2423
22.	Dewing 1660	41.	Dewing 2191	60.	Dewing 2431
23.	Dewing 1669	42.	Dewing 2233	61.	Dewing 2436
24.	Dewing 1699	43.	Dewing 2287	62.	Dewing 2440
25.	Dewing 1703	44.	Dewing 2289		

TABLE B.1 DEWING COINS BY CATALOG NUMBERS

TABLE B.2 is a complete catalog list of the 70 BCMA coins modeled in the coin catalog.

#	Coin Catalog #	#	Coin Catalog #	#	Coin Catalog #
1.	BCMA 1914 6 1	5.	BCMA 1914 6 7	9.	BCMA 1915 98
2.	BCMA 1914 6 2	6.	BCMA 1914 6 8	10.	BCMA 1919 58 1
3.	BCMA 1914 6 3	7.	BCMA 1914 6 12	11.	BCMA 1919 58 2
4.	BCMA 1914 6 6	8.	BCMA 1915 94	12.	BCMA 1919 58 3

#	Coin Catalog #	#	Coin Catalog #	#	Coin Catalog #
13.	BCMA 1919 58 4	33.	BCMA 1919 58 40	53.	BCMA 1920 8 39
14.	BCMA 1919 58 4a	34.	BCMA 1919 58 58	54.	BCMA 1920 8 45
15.	BCMA 1919 58 5	35.	BCMA 1919 58 63	55.	BCMA 1920 8 46
16.	BCMA 1919 58 8	36.	BCMA 1919 58 64	56.	BCMA 1920 8 47
17.	BCMA 1919 58 8b	37.	BCMA 1919 58 65	57.	BCMA 1920 8 49
18.	BCMA 1919 58 8D	38.	BCMA 1919 58 66	58.	BCMA 1920 8 51
19.	BCMA 1919 58 12	39.	BCMA 1919 58 67	59.	BCMA 1920 8 52
20.	BCMA 1919 58 17	40.	BCMA 1919 58 68	60.	BCMA 1920 8 53
21.	BCMA 1919 58 18	41.	BCMA 1919 58 73	61.	BCMA 1920 8 55
22.	BCMA 1919 58 19	42.	BCMA 1919 58 78	62.	BCMA 1923 118 2
23.	BCMA 1919 58 21	43.	BCMA 1919 58 89	63.	BCMA 1923 118 3
24.	BCMA 1919 58 24	44.	BCMA 1919 58 95	64.	BCMA 1923 118 4
25.	BCMA 1919 58 25	45.	BCMA 1919 58 97	65.	BCMA 1923 118 5
26.	BCMA 1919 58 26	46.	BCMA 1919 58 98	66.	BCMA 1923 118 6
27.	BCMA 1919 58 27	47.	BCMA 1919 58 99	67.	BCMA 1923 119 5
28.	BCMA 1919 58 28	48.	BCMA 1919 58 104	68.	BCMA 1923 119 9
29.	BCMA 1919 58 29	49.	BCMA 1920 8 1	69.	BCMA 1958 18 2
30.	BCMA 1919 58 31	50.	BCMA 1920 8 2	70.	BCMA 1958 18 5
31.	BCMA 1919 58 37	51.	BCMA 1920 8 37		
32.	BCMA 1919 58 39a	52.	BCMA 1920 8 38		

TABLE B.2 BCMA COINS BY CATALOG NUMBERS

B.4 LONDON VASE COLLECTION

There are 157 vases modeled in the 3D interactive Greek vase museum described in Chapter 6. These vases are selected from the London collection in Perseus. The vases are identified by the collection and the catalog numbers (e.g. London 1897-7-27-2). Vases with similar catalog numbers tend to share common traits in many cases. At the time of writing the thesis, the current London vase site shows only 156 vases. However, at the time of 3D modeling, there were 157 vases found for the London collection. London K 38 seems to be the extraneous vase. It has no picture associate with it. It is a cup with no texture map. TABLE B.3 is a complete catalog list of the 157 London vases modeled in the vase museum.

#	**Vase Catalog #**	#	**Vase Catalog #**	#	**Vase Catalog #**
1.	London 1897-7-27-2	22.	London 1978-6-15-1	43.	London B 300
2.	London 1899-2-19-1	23.	London 1980-10-29-1	44.	London B 304
3.	London 1899-7-21-4	24.	London A 342	45.	London B 323
4.	London 1899-7-21-5	25.	London B 39	46.	London B 327
5.	London 1901-5-14-1	26.	London B 52	47.	London B 329
6.	London 1903-2-17-1	27.	London B 76	48.	London B 364
7.	London 1912-5-22-1	28.	London B 130	49.	London B 376
8.	London 1914-10-30-1	29.	London B 134	50.	London B 380
9.	London 1917-12-10-1	30.	London B 144	51.	London B 421
10.	London 1917-7-25-2	31.	London B 148	52.	London B 424
11.	London 1920-12-21-1	32.	London B 150	53.	London B 425
12.	London 1923-10-16-10	33.	London B 161	54.	London B 432
13.	London 1923-1-18-1	34.	London B 163	55.	London B 436
14.	London 1928-2-13-1	35.	London B 174	56.	London B 471
15.	London 1928-7-19-3	36.	London B 193	57.	London B 475
16.	London 1929-11-11-1	37.	London B 209	58.	London B 492
17.	London 1947-7-14-18	38.	London B 210	59.	London B 502
18.	London 1961-7-10-1	39.	London B 215	60.	London B 507
19.	London 1964-4-15-1	40.	London B 226	61.	London B 590
20.	London 1971-11-1-1	41.	London B 280	62.	London B 605
21.	London 1977-12-11-9	42.	London B 298	63.	London B 607

#	Vase Catalog #	#	Vase Catalog #	#	Vase Catalog #
64.	London B 610	93.	London E 135	122.	London E 370
65.	London B 658	94.	London E 137	123.	London E 372
66.	London B 668	95.	London E 140	124.	London E 382
67.	London B 674	96.	London E 157	125.	London E 410
68.	London B 675	97.	London E 159	126.	London E 424
69.	London B 693	98.	London E 162	127.	London E 432
70.	London D 2	99.	London E 163	128.	London E 437
71.	London D 4	100.	London E 169	129.	London E 439
72.	London D 6	101.	London E 171	130.	London E 440
73.	London D 7	102.	London E 178	131.	London E 466
74.	London D 13	103.	London E 180	132.	London E 695
75.	London D 20	104.	London E 183	133.	London E 697
76.	London D 48	105.	London E 190	134.	London E 698
77.	London D 51	106.	London E 219	135.	London E 768
78.	London D 58	107.	London E 222	136.	London E 772
79.	London D 60	108.	London E 224	137.	London E 773
80.	London D 61	109.	London E 227	138.	London E 774
81.	London D 71	110.	London E 228	139.	London E 775
82.	London E 3	111.	London E 233	140.	London E 808
83.	London E 23	112.	London E 241	141.	London F 157
84.	London E 34	113.	London E 255	142.	London F 159
85.	London E 38	114.	London E 258	143.	London F 174
86.	London E 41	115.	London E 266	144.	London F 179
87.	London E 44	116.	London E 270	145.	London F 189
88.	London E 65	117.	London E 271	146.	London F 194
89.	London E 66	118.	London E 284	147.	London F 197
90.	London E 104	119.	London E 313	148.	London F 271
91.	London E 129	120.	London E 336	149.	London F 272
92.	London E 134-2	121.	London E 363	150.	London F 284

#	Vase Catalog #	#	Vase Catalog #	#	Vase Catalog #
151.	London F 303	154.	London F 65	157.	London K 38
152.	London F 473	155.	London F 68		
153.	London F 582	156.	London F 90		

TABLE B.3 LONDON VASES BY CATALOG NUMBERS

APPENDIX C

DATA COLLECTION INSTRUMENTS

C.1 INFORMED CONSENT FORM

Project Title: Understanding New Interaction Techniques for the Digital Library

Principal Investigator: Horn-yeu Shiaw
Computer Science Department
Tufts University
hshiaw@cs.tufts.edu
Ph.D. Advisor: Dr. Robert J.K. Jacob
Computer Science Department
Tufts University
Phone: 617-627-3217
jacob@cs.tufts.edu

Your participation in the following experiment is completely voluntary. You are free to withdraw this consent at any time, for any reason, and to request that any data collected be destroyed. If at any time you feel uncomfortable, you may discontinue your participation in the experiment without prejudice to you.

You will be using standard computers and other office equipment, in a conventional office setting in our laboratory.

The purpose of this experiment is to study the use of different interaction techniques for the digital library for viewing a collection of Greek vases.

The procedure is that we will first explain the vase collection and systems to you and then ask you to perform a list of tasks in a limited time frame while using either one of the two systems or both, and ask for your reactions to the different systems.

However, you are free to decline to perform any or all of the tasks on the list.

The whole procedure should take about 1 hour. There will not be any monetary compensation.

You will also be asked to fill in a questionnaire at the end of the experiment.

However, you are free to decline to answer any or all of the questions on the questionnaire.

We will also ask you some questions orally, and you are free to decline to answer if you wish.

Data collected in this experiment will be completely anonymous, and reported only as group averages. We are interested in evaluating the performance of the different programs you are using, not in evaluating the performance of the people themselves.

You consent to release of scientific data resulting from participation in this study to the Principal Investigator for scientific research.

The Principal Investigator assures your anonymity. The record of this experiment becomes part of the experimental results and is protected as a confidential document. This record will only be available to the researchers involved with this study.

Other staff may be authorized by the Tufts board to review the record for administrative purposes or for monitoring the quality of subject care.

If you have any questions, at any point during the experiment, the experimenter will gladly answer them.

Please read the following and sign on the lines below:

"I, the undersigned, have read and understood the explanations of the following research project and voluntarily consent to my participation in it. I understand that my responses will remain confidential and that I may terminate my participation at any time.

In the unlikely event of physical injury resulting from participation in this research, I understand that medical treatment will be available from the Tufts University Health Services at 627-3350, including first aid emergency treatment and follow-up care as needed, and that my insurance carrier may be billed for the cost of such treatment. However, no compensation can be provided for medical care apart from the foregoing. I further understand that making such medical treatment available; or providing it, does not imply that such injury is the Investigator's fault. I also understand that by my participation in this study I am not waiving any of my legal rights.

I understand that I may also contact the Chairman of the Committee on the IRB, Tufts University, if I feel I have been treated unfairly as a subject.

I confirm that I have passed my eighteenth birthday, the required minimum age necessary to take part in an adult research study."

Signature: ____

Name: ____

Date: ____

Location: Tufts University, Halligan Hall

C.2 TASK QUESTIONS

TASK INSTRUCTIONS

SUBJECT ID: ________

TASK ________

For this task, you are to look at an entire collection of London vases and answer the task questions and associated survey.

You do not have to study each vase in detail; just get a general idea of the whole collection.

You are to concentrate in the appearance of the vases, such as colors, shapes, handles, decorations and artworks.

You are not required to count or read any documentation, descriptions, or writings about the vases beyond what you believe are necessary to perform this task.

You are not required to have any prior knowledge of the vases or archaeology.

Please circle your answers.
Work as quickly as you can, without making errors.

Warm up Question

Sample Task Question:

What is the difference between vase "London B 436" and "London A 342"?

a.) London B 436 has two handles and London A 342 has one handle.

b.) London B 436 has an oval opening and A 342 has an arched neck.

c.) London B 436 is brown/black and A 342 is beige/taupe.

d.) All of the above.

e.) Cannot be determined.

Task Question #1:

What is the most distinctive difference between vase "London B 436" and other vases in this collection?

a.) London B 436 is completely different in color scheme.

b.) London B 436 has unique number(s) of handle(s).

c.) London B 436 has very ornate artwork around the brim.

d.) London B 436 is placed on its side and all other vases are placed upright.

e.) Cannot be determined.

Task Question #2:

Which WARE of vases is mostly made throughout the period 300 – 700 B.C.?

a.) Red figures

b.) Geometric

c.) Black figures

d.) Cycladic

e.) Cannot be determined

Task Question #3:

Which statement below is NOT true?

a.) Many vases have a base.

b.) Many vases have a spout.

c.) Many vases have a stem.

d.) Many vases have a cover.

e.) Cannot be determined.

Task Question #4:

What is the number of handles on the vases?

a.) None

b.) One

c.) Two

d.) All of the above

e.) Cannot be determined

Task Question #5:

Approximately how many vases have NO image of decoration or artwork?

a.) None 0%

b.) Less than 25%

c.) Greater than 25%

d.) All 100%

e.) Cannot be determined

Task Question #6:

What will you NOT find as a theme for decoration and artwork on the vases?

a.) People and animals

b.) Scenes from Greek mythology

c.) Texts and inscriptions

d.) Geometric designs

e.) Cannot be determined

Task Question #7:

Which color will you NOT find in this collection?

a.) Blue

b.) Black

c.) Gold

d.) Brown

e.) Cannot be determined

Task Question #8:

What can NOT be found in this collection?

a.) Lids and covers

b.) Plates and dishes

c.) Cups and bowls

d.) Pots and Pans

e.) Cannot be determined

Task Question #9:

Which best describes the size of the vases in this collection?

a.) Most vases are all the same size

b.) Most vases are similar in size

c.) Most vases are very different in size

d.) Most vases are gigantic in size

e.) Cannot be determined

Task Question #10:

What are the two most common shapes in this collection?

a.)  **Amphora**  **Stamnos**

b.)  Lebes  Loutroforus

c.)  Kantharos  Volute Krater

d.) None of the above

e.) Cannot be determined

C.3 TASK SURVEY

Please circle your opinion on a scale of 1 (Strongly Agree) – 7 (Strong DisAgree):

	In answering this question, I believe …	Strongly Agree	Agree	Slightly Agree	Neutral	Slightly Disagree	Disagree	Strongly Disagree
S.1.1	This question was very easy.	1	2	3	4	5	6	7
S.1.2	I know I got this question right.	1	2	3	4	5	6	7
S.1.3	I can recall this information if asked again.	1	2	3	4	5	6	7
S.1.4	The system is very effective for this question.	1	2	3	4	5	6	7
S.1.5	My answer was based on looking at many vases.	1	2	3	4	5	6	7
S.1.6	I did NOT guess at random.	1	2	3	4	5	6	7

The following two survey questions are not applicable to the warm-up or the first question, but will appear in subsequent surveys:

S.1.7	I was able to answer this question using carry-over knowledge from the previous question(s).	1	2	3	4	5	6	7
S.1.8	This question was easier than the previous question.	1	2	3	4	5	6	7

C.4 QUESTIONNAIRE

C.4.1 PART A : SURVEY

PART A

Please indicate your agreement or disagreement by circling a number on the scale provided.

In looking at the collection and performing the task, I would say …		Strongly Agree	Agree	Slightly Agree	Neutral	Slightly Disagree	Disagree	Strongly Disagree
A.1	I am new to the collection before this experiment.	1	2	3	4	5	6	7
A.2	I have learned a lot about the collection.	1	2	3	4	5	6	7
A.3	I like the system very much.	1	2	3	4	5	6	7
A.4	It is an enjoyable learning experience.	1	2	3	4	5	6	7
A.5	There is sufficient time to answer the questions.	1	2	3	4	5	6	7
A.6	No counting is required to answer the questions.	1	2	3	4	5	6	7
A.7	No reading is required to answer the questions.	1	2	3	4	5	6	7
A.8	I can compare two vases easily.	1	2	3	4	5	6	7
A.9	I based my answers by comparing enough vases.	1	2	3	4	5	6	7
A.10	3D graphical models of the vases are good to have.	1	2	3	4	5	6	7
A.11	I can explain the vase collections to others.	1	2	3	4	5	6	7
A.12	I can compare two groups of vases conveniently.	1	2	3	4	5	6	7
A.13	I can quickly pick out unique-looking vases.	1	2	3	4	5	6	7
A.14	I can compare a vase with many other vases.	1	2	3	4	5	6	7

A.15	I have looked at nearly all the vases.	1	2	3	4	5	6	7
A.16	I have a good idea of what most vases look like.	1	2	3	4	5	6	7
A.17	I can identify most vases if I see them again.	1	2	3	4	5	6	7
A.18	I can recognize the vases in a real museum.	1	2	3	4	5	6	7
A.19	I think the vases are visually appealing.	1	2	3	4	5	6	7
A.20	Greek vases can be different from today's vases.	1	2	3	4	5	6	7

C.4.2 PART B : COMMENTS

PART B

B.1 Describe how you viewed and navigated through the collection.

B.2 Describe the process of how you arrived at your answers.

B.3 Summarize what you have learned about the vases in this collection.

B.4 Why do you think you needed to guess at random? (Skip if you did not guess at random)

B.5 STRONG points of the system, including what you LIKE the most.

B.6 WEAK points of the system, including what you DISLIKE the most?

B.7 What improvements to the system would you like to see implemented?

B.8 Please give us any comments or suggestions you have on this experiment:

C.4.3 PART C : QUIZ

PART C *(Answers are in bold)*

Please circle *mostly* TRUE or *mostly* FALSE to each question.

You may not look at the vases again. Try your best to answer from memory.

	I would say …	***Mostly***	***Mostly***
C.1	Red figure vases were made throughout the periods.	**TRUE**	FALSE
C.2	Black figure vases were most popular in 200’s B.C.	TRUE	**FALSE**
C.3	There were more red figure vases than black figure vases.	**TRUE**	FALSE
C.4	Color of red figure vases was different shades of red.	TRUE	**FALSE**
C.5	Color of black figure vases was different shades of black.	TRUE	**FALSE**
C.6	Color of white ground vases was different shades of white.	**TRUE**	FALSE
C.7	There were purple and blue vases.	TRUE	**FALSE**
C.8	Most vases were brown and black.	**TRUE**	FALSE
C.9	Vase London B 436 lay on its side.	**TRUE**	FALSE
C.10	Vase London A 342 had an arched neck	**TRUE**	FALSE
C.11	Quite a few vases had text and writings.	TRUE	**FALSE**
C.12	People and animals were depicted on most vases.	**TRUE**	FALSE
C.13	Flowers and plants were depicted on many vases.	TRUE	**FALSE**
C.14	Whole-body shapes of many vases could not be determined.	TRUE	**FALSE**
C.15	Nearly all vases were red figure vases.	TRUE	**FALSE**
C.16	Images of artwork were shown on all vases.	TRUE	**FALSE**

C.17	Many vases had openings facing upward.	**TRUE**	FALSE
C.18	No vase had a cover.	TRUE	**FALSE**
C.19	Some vases had no handle	**TRUE**	FALSE
C.20	Many vases had one handle.	**TRUE**	FALSE
C.21	Many vases had two handles.	**TRUE**	FALSE
C.22	Some vases had three handles.	TRUE	**FALSE**
C.23	Vases made in 300 – 400 B.C had simpler shapes.	TRUE	**FALSE**
C.24	Most vases were made in 450 – 550 B.C.	**TRUE**	FALSE
C.25	Vases made in 600 – 700 B.C. were much fancier.	TRUE	**FALSE**
C.26	As time progressed, the vases were more complex in shape.	TRUE	**FALSE**
C.27	As technology matured, vases became much more colorful.	TRUE	**FALSE**
C.28	Vases in the collection were very different in color.	TRUE	**FALSE**
C.29	Most vases had a base.	**TRUE**	FALSE
C.30	Some vases had a stem.	**TRUE**	FALSE
C.31	The ‘pan’ was a popular shape in this collection.	TRUE	**FALSE**
C.32	Many vases had a spout.	TRUE	**FALSE**
C.33	A typical vase had an opening, round body and two handles.	**TRUE**	FALSE
C.34	Greek vases could not include lids.	TRUE	**FALSE**
C.35	Greek vases could not include cups.	TRUE	**FALSE**
C.36	Greek vases could not include plates.	TRUE	**FALSE**
C.37	Greeks decorated most of their vases with artwork.	**TRUE**	FALSE
C.38	Many vases appeared to be similar in size.	**TRUE**	FALSE

C.39	Shapes of many of the vases cannot be determined.	TRUE	**FALSE**
C.40	Some vases had more information than others.	**TRUE**	FALSE
C.41	Vases were distinguished by shape, ware and year.	**TRUE**	FALSE
C.42	There were only a few different shapes of vases.	**TRUE**	FALSE
C.43	It was unclear what the vases were used for.	**TRUE**	FALSE
C.44	Catalog numbers were used to identify the vases.	**TRUE**	FALSE
C.45	Painters of most vases were known.	TRUE	**FALSE**
C.46	Most artwork on many vases depicted mythological scenes.	**TRUE**	FALSE
C.47	Greek vases were mostly used to hold and display flowers.	TRUE	**FALSE**
C.48	Many vases were made in unknown regions.	**TRUE**	FALSE
C.49	Some vases were black and white.	**TRUE**	FALSE
C.50	Other Greek vase collections were similar to this one.	**TRUE**	FALSE

APPENDIX D

DATA COLLECTED

There were 20 subjects, 12 male and 8 female, divided into two groups:

V – The group of 10 subjects who used the 3D interactive virtual museum.

There were 6 male (M) and 4 female (F) subjects in V.

Each subject was identified by a code: V[01..10]-[M | F]-S[01..20]

e.g. V1-M-S1 was the 1st subject, male, in V and the 1st in the experiment.

P – The group of 10 subjects who used the Perseus Digital Library

There were 6 male (M) and 4 female (F) subjects in P.

Each subject was identified by a code: P[01..10]-[M | F]-S[01..20]

e.g. P3-F-S6 was the 3rd subject, female, in P and the 6th in the experiment.

A Likert Scale of 1 (Strongly Agree) – 7 (Strongly Disagree) was used.

T-test p-value used "same equal variance" and "two-tailed distribution" and the p-value must be < 0.05 to be significant.

D.1 USER OBSERVATIONS

3D Interactive Vase Museum	
Subjects	Observation Notes
V01-M-S01	Not used to keypad on laptop; Impatient about navigation or changing view ports; Found anitour, and used it for most questions; Said did not like to use anitour, but seemed to be the best way to look at the vases.
V02-M-S02	Not used to keypad on laptop, but got used to it soon; Stayed on a few view ports and familiar vase scenes.
V03-M-S04	Very inefficient user; Could be new to computers; Blamed himself constantly for not doing well; Used navigation as much as view ports.
V04-F-S08	Based her answers on seeing just a few vases, so was able to go really fast on some questions; Spent a lot of time navigating, unlike others who tend to use the view ports more.
V05-F-S11	Really fast; Based answers on 2 – 3 view ports so this subject was very quick in her responses; Did not know what "wares" meant even after explanation.
V06-M-S12	Found anitour and stayed on it for the rest of questions; Used scatter plot quite well; Tried to click on the vase dots on scatter plot to go to them.
V07-M-S13	Really efficient; Picked a few view ports for each new question and looked at the background vases.
V08-F-S17	View port user; Read questions, picked a few view ports, thought for a while, repeated the process one or more times before arriving at answers.
V09-F-S19	Used navigation a lot; Sometimes had to wait a long time for her to 'walk' to the vase; Initially tried to see the back of the vases; Gave that up after realizing that the browser was not very maneuverable.
V10-M-S20	Talked to me in HCI 'think out loud'; Liked the scatter plot, read it really well; Tried to get a good view of the vases by navigation, stalling; Liked to go into the 2D HTML pages and rotate the vases.

Tufts Perseus Digital Library	
Subjects	Observation Notes
P01-M-S03	Browsed the list; Arrowed up and down the list; Looked at list; A well-behaved user who did what was supposed to be done.
P02-M-S05	Used search to count; Actually counted question #10; Wanted to use google.com to answer the questions (rather than using Perseus); Savvy user; Searched HTMLs in notepads, which was totally unnecessary.
P03-F-S06	Did not browse in many situations; Guessed by seeing a few; Preferred to read title page; Said users are placed on the spot, detached from HCI designers; Said "effective" was not a fair question; Said would do better if given enough time to read; More vocal; preferred not to write down comments; Said it was English that got in the way.
P04-F-S07	Tried to answer without using Perseus at all (thought this was an exam); Tried to type questions in the search engine; Went to different collections; Considered "Text and Inscriptions" to mean HTML writings.
P05-F-S09	Repeatedly got out of the London collection; Tried to answer by search; Wanted to type in questions as is into the searches; Insisted on doing it even though Perseus wouldn't take.
P06-M-S10	Shortcut-key user; Jammed screens with lots of keystrokes; Very impatient about waiting for the HTML pages to load; Said Perseus was very slow; Server errors often happened.
P07-M-S14	Got 9 out of 10 questions right in a short time; Did not ask questions; Fast; Seemed to know some of the answers without even looking at the system.
P08-M-S15	Talked a lot; 'think out loud' turned into non-stop dictation; Cleared out menu bar and tool bar; Right click to open windows; Opened too many windows; Perseus did not like that, so little or no image was loaded; By question #7, yawning and impatient, did not talk much any more.
P09-M-S16	Well behaved; Browsed, viewed the HTML pages, and answered the questions; Went really smoothly.
P10-F-S18	Did not browse, or see system at all on many questions; Seemed to answer on intuition or based on the first few seen.

D.2 TASK QUESTIONS

3D Interactive Vase Museum			
Subjects	Warm up (1)	Task (10)	Time (in min)
V01-M-S01	1	8	15.0170
V02-M-S02	1	8	18.6300
V03-M-S04	1	7	20.3650
V04-F-S08	1	6	19.7150
V05-F-S11	1	5	07.1070
V06-M-S12	1	8	09.2303
V07-M-S13	1	7	08.8873
V08-F-S17	1	7	12.2370
V09-F-S19	0	8	10.2920
V10-M-S20	1	8	15.3950
N	10	10	10
Sum	9	72	136.88
Mean	0.9	7.2	13.688
Min	1	8	20.365
Max	0	5	7.107
Variance	0.1000	1.0667	23.3196
Stdev	0.3162	1.0328	4.8290
T-test p value	0.1934	0.0074	0.0001

Tufts Perseus Digital Library			
Subjects	Warm up (1)	Task (10)	Time (in min)
P01-M-S03	0	6	32.5970
P02-M-S05	1	6	27.9380
P03-F-S06	1	4	22.9310
P04-F-S07	0	4	37.8890
P05-F-S09	1	4	52.2890
P06-M-S10	1	6	64.5030
P07-M-S14	1	9	22.5560
P08-M-S15	0	4	50.9860
P09-M-S16	1	6	31.9090
P10-F-S18	0	5	26.6580
N	10	10	10
Sum	6	54	370.256
Mean	0.6	5.4	37.0256
Min	1	9	64.503
Max	0	4	22.556
Variance	0.2667	2.4889	203.0564
Stddev	0.5164	1.5776	14.2498

D.3 TASK SURVEY

Task Survey – Summary of Means													
20 Survey Questions 1 –7 (lower #s better)			20 Subjects (10 in V and 10 in P)										Grand Means
			1	2	3	4	5	6	7	8	9	10	
1	SENSE OF QUESTION EASINESS	V	4.7	3.5	3.7	3.3	3.6	2.7	3.2	3.2	3.5	5.2	3.7
		P	4.4	3.5	4.6	4.4	3.9	2.7	3.5	3.6	2.3	4.1	3.7
2	SENSE OF ANSWER ACCURACY	V	4.8	3.5	3.2	2.9	3.9	2.8	3.1	2.7	2.9	4.5	3.4
		P	4.2	2.5	4.4	4.0	3.3	2.8	3.9	3.4	2.1	4.2	3.5
3	SENSE OF ABILITY TO RECALL	V	3.4	2.1	2.9	2.0	3.4	2.5	3.2	3.2	3.7	4.3	3.1
		P	3.0	2.1	1.0	2.5	2.7	2.2	2.9	3.5	1.0	4.4	2.5
4	SENSE OF SYSTEM EFFECTIVENESS	V	4.8	3.9	2.3	2.5	3.9	2.4	4.1	3.7	4.1	3.8	3.6
		P	3.9	4.7	4.0	4.6	3.7	2.2	4.2	3.5	3.5	5.7	4.0
5	SENSE OF ENOUGH VASE SAMPLING	V	2.4	4.7	1.7	1.6	2.6	3.1	3.2	1.3	2.2	1.8	2.5
		P	2.6	2.1	4.0	4.0	3.3	2.8	3.4	3.0	6.4	4.2	3.6
6	SENSE OF NO GUESSING IN ANSWER	V	2.9	1.9	2.1	1.8	1.4	2.7	3.5	1.7	2.7	3.6	2.4
		P	3.5	1.7	2.6	3.4	3.7	2.7	3.1	2.4	3.1	4.2	3.0
7	SENSE OF CARRY-OVER KNOWLEDGE	V	3.1	3.6	5.4	1.8	2.2	3.9	2.6	1.4	2.2	2.8	2.9
		P	2.8	2.7	2.9	4.1	3.1	2.3	3.6	2.7	1.8	3.4	2.9
8	SENSE OF ACCUMULATED LEARNING	V	4.2	3.9	5.7	3.7	4.3	3.3	3.3	3.8	4.8	4.1	4.1
		P	3.7	3.8	3.6	4.1	3.2	3.9	4.1	3.2	3.0	4.2	3.7

D.4 QUESTIONNAIRE

D.4.1 PART A : SURVEY

Questionnaire Survey – Summary of Means

Survey Questions 1–7 (lower #s better)			20 Subjects (10 in V and 10 in P)										Grand Means
			1	2	3	4	5	6	7	8	9	10	
1	NEW TO THIS COLLECTION	V	1	1	1	1	1	2	1	1	7	1	1.7
		P	5	1	1	1	2	1	5	3	2	1	2.2
2	LEARNED A LOT ABOUT THIS COLLECTION	V	5	5	2	2	5	3	3	4	5	4	3.8
		P	3	5	4	7	2	3	5	2	3	6	4.0
3	LIKE SYSTEM VERY MUCH	V	7	5	1	2	3	3	5	3	4	3	3.6
		P	3	6	4	4	4	4	7	2	3	7	4.4
4	ENJOYABLE LEARNING EXPERIENCE	V	7	5	1	2	2	2	4	3	3	4	3.3
		P	3	2	4	3	3	2	6	2	2	6	3.3
5	SUFFICIENT TIME TO ANSWER	V	1	1	1	2	2	2	2	3	2	2	1.8
		P	5	2	3	5	2	2	1	2	1	6	2.9
6	NO COUNTING RQUIRED	V	1	3	1	2	4	5	2	1	2	3	2.4
		P	5	6	1	6	2	6	3	6	5	4	4.4
7	NO READING REQUIRED	V	1	1	1	1	2	3	5	2	2	5	2.3
		P	5	6	1	2	4	2	5	6	5	4	4.0
8	COMPARE TWO VASES EASILY	V	5	2	1	2	2	6	2	2	3	6	3.1
		P	6	6	2	4	3	3	6	4	4	4	4.2
9	COMPARED	V	2	4	1	2	3	3	3	1	3	1	2.3

	ENOUGH VASES	P	3	3	4	5	2	2	4	2	1	6	3.2
10	GOOD TO HAVE 3D MODELS OF VASES	V	1	1	1	1	1	3	1	3	2	2	1.6
		P	2	5	4	2	2	2	1	5	1	2	2.6
11	CAN EXPLAIN COLLECTION TO OTHERS	V	3	6	3	2	3	4	3	3	3	7	3.7
		P	2	4	4	6	2	2	3	2	3	4	3.2
12	COMPARE TWO GROUPS OF VASES EASILY	V	7	5	2	2	5	5	5	3	3	6	4.3
		P	6	7	4	6	2	4	4	4	4	6	4.7
13	PICK OUT UNIQUE VASES	V	7	3	2	2	2	5	2	3	4	6	3.6
		P	5	7	4	6	3	5	7	5	4	6	5.2
14	COMPARE A VASE WITH MANY OTHERS	V	6	6	3	5	2	5	5	3	4	6	4.5
		P	3	6	4	4	3	1	3	3	3	6	3.6
15	LOOKED AT NEARLY ALL VASES	V	7	6	3	2	3	5	5	4	4	5	4.4
		P	5	6	5	5	4	3	6	2	6	7	4.9
16	KNOW WHAT MOST VASES LOOK LIKE	V	6	2	3	2	2	3	3	2	5	5	3.3
		P	4	3	5	3	4	2	2	1	2	6	3.2
17	IDENTIFY VASES IF SEE THEM AGAIN	V	7	2	4	3	2	5	3	3	6	6	4.1
		P	5	6	5	5	3	2	5	3	3	6	4.3
18	RECOGNIZE VASE IN A REAL MUSEUM	V	4	4	3	2	2	5	5	3	6	6	4.0
		P	3	6	4	6	2	4	5	4	2	5	4.1
19	VASES ARE VISUALLY APPEALING	V	3	2	2	2	2	3	4	2	4	4	2.8
		P	3	3	1	2	2	4	2	2	1	4	2.4
20	VASES LOOK DIFFERENT FROM TODAY'S	V	4	1	4	2	2	4	3	3	4	3	3.0
		P	2	3	2	1	2	2	1	1	3	5	2.2

D.4.2 PART B : COMMENTS

3D Interactive Vase Museum	
#	Subject : V01-M-S01
B.1	Too often, I went to the "anitour" view as it seemed to be the only overview of the whole collection, but I didn't like it. For some questions, I went to the look-down-from-ceiling (scatter plot) view. I never used any of the "study" options because the room was too big to "study."
B.2	Frustratingly, I would have liked to "walk" around the room, taking a closer look occasionally, but all I could do was take a far view. Walking was so slow & unnatural, it was useless.
B.3	There were many kinds, usually open at the top. Many were shaped like [drew something like a light bulb] and many were more bowl-like. Most had artwork, some had one handle & a spout; some had 2 handles, some none. There were made between 300 – 700 B.C.
B.4	Could not get a good overall sense of what was there.
B.5	Visual aspect – able to see the vases almost as if I were there.
B.6	Speed was so slow and it was frustrating. Hard enough to use a new system with a mouse. Notebook made it even worse.
B.7	A real-time (preferably one-hand) control that lets you wander around the room, look at vases, etc.
B.8	A lot of the questions were ambiguous or confusing.
#	Subject : V02-M-S02
B.1	Initially browsed by each individual vase, but then just looked at collections to get an idea of what's out there.
B.2	Conclusion was made early on that vases are very similar in color and design. I initially thought there were 4 – 5 different types of vases, but then I needed to check again for additional ones. I checked every answer to answer best I can, but some questions required no browsing.

B.3	– 6 – 7 different types – brown, gold, black, pictures on them – some are completely white – scenes from mythology, humans & animals
B.4	Mostly I didn't guess at random.
B.5	– 3D vases are excellent idea – colors are easy to distinguish
B.6	– navigation – items should be grouped differently – blurry pictures, hard to say what is on them – descriptions of each individual vase is not clear enough – Slow – Hard to distinguish size of vases
B.7	– different navigation – when individual vase is selected, have a photo of it appear as well. – possibly more charts with relevant info, such as initial chart with wares
B.8	– didn't like the laptop mouse, would prefer regular
#	Subject : V03-M-S04
B.1	Using the mouse, pointer and directional keys I navigated through the collection. I became more adapted as I proceeded through the test.
B.2	I answered @ my answers by comparison, viewing and studying the vases to derive my answers.
B.3	That the collection was diverse, and spanned many centuries. That they were able to give me an accurate representation of the actual vase.
B.4	
B.5	That I was able to thoroughly view the collection before giving my answers.
B.6	Difficulty in learning the system at the beginning. The poor response time @ the mist of testing.
B.7	Brief written instructions & graphics to explain use + navigation
B.8	It was very insightful.

#	Subject : V04-F-S08
B.1	At first, I selected a vase from menu. Then I go around the vase to see other vases. When I found I can select to view all the vases of B.C. 300 or B.C. 400, I choose the group view. Then go to the vase which I am interested in.
B.2	1. View a group of vases to find the common properties of them. Then go to another group of vases to find their common properties. 2. If I need to compare two vases, look at them separately.
B.3	1. The colors of most vases are black, brown, white. 2. No text on the vases 3. The shapes of most vases are [drew a light bulb shape] 4. Most vases have no lips 5. A few pans 6. Most white vases are similar – Some vases are big
B.4	
B.5	1. I can go around vases, view them at different viewpoint, just like in a museum. 2. I can find a vase quickly.
B.6	Slow speed.
B.7	If there is a map which indicates where I am, it will be better.
B.8	1. If put the group view (vases before B.C. 300 etc.) on the top of the menu, it will be convenient. 2. Group the vases according to date in the long menu.
#	Subject : V05-F-S11
B.1	Using the next, prev arrows; Using the view pop up menu – not convenient; I couldn't select an overview of many vases vs. specific vase.
B.2	Tried to view as many vases as I could.
B.3	Shapes, variety, colors, dimensions compare to each other
B.4	
B.5	3D image of a vase; View of many vases at a time

B.6	Selecting the view; Slow
B.7	Faster, better way to select the view, better way to navigate.
B.8	
#	Subject : V06-M-S12
B.1	I mostly right-clicked and picked a few viewpoints at random.
B.2	I viewed several vases and made an educated guess.
B.3	I haven't learned specifics – the main fact I recall is that most of the vases in the collection were red vases.
B.4	My patience could not last to examine every vase.
B.5	I like the visuals and the organization of the museum.
B.6	It was difficult to compare vases – perhaps two windows would work better? And speed of navigation was difficult.
B.7	Multiple viewpoints; Faster speed
B.8	Explain that the entire collection of vases is the "London collection"
#	Subject : V07-M-S13
B.1	When a question asked about the group of vases, I would look at the ones in the background to get a general idea.
B.2	Most questions asked about the collection, so I would think about all those I had already seen. If I couldn't answer it, I would randomly pick an element from the list, and look at those in the background.
B.3	I learned a little about the shapes & sizes, including the different handle configurations. I also learned about when time period in which these vases were made was.
B.4	Because the questions were asking about the entire group of vases and I hadn't seen them all. Also when the info couldn't be determined from the picture, and more in-depth reading/research would have been required.

B.5	Different zoom levels allow you to get a general idea about the group. Zooming in gives more details. Conceptually this is intuitive.
B.6	Slow computer system meant I was hesitant to move in the environment or make changes. Also, when close to a vase, the pixilation/distortion made discerning details difficult. This meant questions about things like what was pictured were impossible.
B.7	Better “fly-through” navigation. A way for getting textual info w/out loading another page.
B.8	
#	Subject : V08-F-S17
B.1	Sometimes Randomly; Sometimes following the time.
B.2	Find and watch lots of vases in different years.
B.3	Vases are beautiful. There have some rules to arrange them.
B.4	
B.5	Can find the year of the vase very easy.
B.6	The imaging of the vases are not good enough.
B.7	I’d like to see very beautiful pictures with many details on it.
B.8	Make more attractive pictures. This system can only turn around from left to right, not up to down.
#	Subject : S09-F-S19
B.1	
B.2	Tried to look around @ as many as I could.
B.3	They are group by age/material
B.4	Machine too slow
B.5	Being able to look around

B.6	To slow especially moving forward.
B.7	Faster
B.8	
#	Subject : V10-M-S20
B.1	I choose a particular vase randomly from the list, hoping that by viewing it, I would also see many vases in the background. Navigating via the keyboard was painfully slow; I abandoned that quickly. The exception to this navigation technique was the question about 2 specific vases; there I obviously did not choose randomly from the list.
B.2	I navigated as described above, looking for criteria sought by the question (vase color, for example). I usually chose about 4 – 6 vases from the list before answering. However, each vase had 3 – 8 other vases in the background, so I viewed 12 – 48 vases for each question. Since I thought it would be inefficient to view all 150 vases, answers were often based on estimates (but not random guesses, usually).
B.3	1. There are few/no blue vases. 2. Only one vase is not standing upright 3. There are few/no cups & bowls in the collection. 4. The vase laying on its side had an oval/circular opening 5. Most vases in the collection are Red Figures. 6. All vases were made between 300 and 700 B.C. 7. There are few/no vases whose design is based on text or inscriptions. 8. Most vases have stems and bases.
B.4	It would have been too time-consuming to individually examine each vase. (This applied only to question about the design/decoration of vases, where text/inscription was one of the answers. It was hard to see enough detail in the main room/window to determine if text was written on the vase. The decoration was too pixelated.)
B.5	– felt like I was in a museum, not just a web page. – walking through the room, I could usually see the timeline in the background, so I know the time period of the vase in front of me. – it was easy to get a vague, general sense of the collection b/c of the 3D room.
B.6	– keyboard navigation has a horrible latency; redraws are too slow – up-close views of vases are too pixelated to allow for detailed analysis

	– It wasn't clear if I could hover/fly over the vases, starting from the overview perspective
B.7	It sounds trivial, but just resolve the dislikes from question B.6. I guess it might be slick to have audio feedback on each vase, like a guided museum tour.
B.8	It wasn't clear whether the "I viewed many vases" question applied to that last task, or overall in the experiment.

Tufts Perseus Digital Library	
#	Subject : P01-M-S03
B.1	Generally I browsed the list from top to bottom looking for hints at particular features I was trying to investigate. Then I would click on an item to view them in detail, generally leading to one or two new browsing sessions based on information learned. Generally I had to return to already viewed vases to remind myself of some things already seen.
B.2	Sometimes by looking for examples to rule out an answer (this was not always as effective as I have hoped due to time limitations.) Sometimes by browsing collections list and trying to get a sense of amounts of a particular feature (felt more effective). * clicking for more detail was sometimes necessary.
B.3	There seem to be several major style types, i.e. attic red, attic black, geometric. And several related shapes (amphora → other names escape me); Sensed that most were from earlier in the time period (i.e. 400 or so and earlier) i.e. not many pass 700. Also, most vases have a story.
B.4	Time, I suspect given enough time, I could arrive at sure answers. Often had to choose between two answers, which I was sure of but not positive.
B.5	Pictures. Summaries, descriptive text (i.e. height since pictures alone do not tell this) description of meaning of pictures at vase → good length of short description in list → helped me look (guess at) details.
B.6	Inconsistency → sometimes due to images loading slowly. Also missing information (i.e. no height). Slow click through to detail. I would have like to be able to selectively choose detail if I could get faster loading.
B.7	See last question (more levels of detail). Also, more summary info on list page. Additionally, dates more prominently displayed would give me a greater sense of

	historical context immediately upon viewing details of a vase.
B.8	Most questions seemed difficult, some more that were really easy (I only felt one was like this). Might have made others seem easier. Whole experience made me feel web is less effective than I thought! (but it seemed better as time went on.)
#	Subject : P02-M-S05
B.1	I used a combination of browser-based text searches & multiple windows drawn from internal links. The system appeared to have some issues with concurrent requests.
B.2	For answers requiring elimination, I looked for examples or counter-examples to the proposed assertions. For those requiring generalization, I tried to view data on a "good" sample set by viewing objects well-distributed along the index page.
B.3	Most of them appear to be stored in London.
B.4	
B.5	I wasn't able to use the display filters, but they seemed quite useful.
B.6	* physical measurements didn't use consistent units. * images were not consistently displayed on the first level – even when available in the system & accessible by further inquiry. * standardization of descriptions is somewhat lacking. * my educated guess is that the speed of results delivery can be significantly improved.
B.7	Approximation of speed * (1) filter at all levels for (2) better standardized data * better –standardized presentation of all available data * specifically, standardized descriptions of shape & features of the vases would help * configurable search would certainly be helpful.
B.8	["Did you view many vases …?" <u>this time</u>, or across all questions?]
#	Subject : P03-F-S06
B.1	
B.2	

B.3	
B.4	The machine is too slow when showing a picture.
B.5	
B.6	
B.7	
B.8	Should put the user in context. Give a brief introduction of the system. Otherwise, it's unfair for both the user and the system designer – a more informed user will be a happier user.
#	Subject : P04-F-S07
B.1	Kind of slow in navigation. I was instructed to only use the "back" button (of IE) to do my browsing.
B.2	Generally, I browse a few vases, then got my answers based on rough probability.
B.3	I have an impression on the decoration of Greek ancient vases. They usually come from mythology, war, or their life. The vases are beautiful, though I didn't learn much during 30'.
B.4	Some question needs guess 'cos it's unlikely possible to find quickly like "which color did not on the vases."
B.5	The decoration description. It's the most attractive feature of the system.
B.6	The category looks weak. I didn't find any other category except "London, Boston, .." I had to browse the vases one by one.
B.7	To make more category, make the searching and sorting powerful.
B.8	I feel frustrated when I can't determine most question. The experiment may include some straight foreward questions to make the testee a little more comfortable.
#	Subject : P05-F-S09
B.1	I was given a whole list of London, British museum collection. As to every question I was confronted, I needed to browse the link one by one by clicking on my mouse.

B.2	I used the search engine on the website by entering some keyword. For instance, I entered "London B 436" to view its image and detailed description. It was easier to get some info like color or how many handles the vase has. For complicated question like how many handles most of the vases have. I have to compare five or six of them.
B.3	I learned the main trend of Greek vase ant. Especially the basic color scheme, shape and year. As far as the color is concerned, most of them are black and brown or gold, shape can be very various in details such as arched neck, spout on long stem. The dimension info is kinda necessary to get a whole picture of how actual size of each master piece. There are about similar sizes of 20 – 40 cm.
B.4	Because I had to draw my conclusion based on comparison among a whole lot of vases, which is sorta arbitrary and subjective. I did not guess totally at random. However, due to my poor knowledge of vase and the limited amount of info I can find online, it is not very convincing for myself to get a very positive answer.
B.5	The vivid pictures of each vase is definitely the strongest point of the system, as well as the description underneath.
B.6	Searching throughout the whole system can be very troublesome and painful due to the nature of the questions I was given. For instance, I was asked to find out the percentage of the vases without decoration and artwork. It is supposed to exist some informative articles containing these kind of info which should be posted somewhere on the web. So we can browse it and get the piece of info easily and quickly.
B.7	The system should contain some art criticism articles so that the readers are able to get a more general picture of the specific topic. The info of each item should be more consistent, namely, using the same units (height in cm or in m). Some item are short of dimension info on date of production and such.
B.8	To be more tolerant and sensitive to users' searching keyword. The GUI should be more user friendly in other words.
#	Subject : P06-M-S10
B.1	Question-based navigation by comparing many vases to either identify the correct answer or rule out the incorrect. And most of time I say the picture only. I randomly choose a vase to take a look.
B.2	As B.1

B.3	I saw many vases all of them have complex decorations and symbols, but rare of them have text or inscriptions. They like to describe lives using pictures.
B.4	The network speed is slow. I waited for 5 minutes without response.
B.5	Can provide an overview to the Greek vase. I didn't read descriptions of each vase, but I have a whole picture.
B.6	Slow response.
B.7	I think 3-D simulation will be an improvement provided network is fast enough.
B.8	
#	Subject : P07-M-S14
B.1	I navigated through the collection by clicking on the links for the piece I wanted to view. If I couldn't make out the vase from the thumbnail, I clicked on it to get a bigger view.
B.2	My answers were a combination of loading as selections from the collection and making educated guesses based on the info I saw.
B.3	There is a wide range of vases. Many are traditional amphoras while others are less common style, most are in either red figure or black figure.
B.4	It would be necessary because the site isn't that easy to navigate. I couldn't sort the list on attributes of the vases (i.e. wares) if I wanted to.
B.5	Thumbnails, when they existed, having attribute info on the main page.
B.6	Sorting. Can't sort list on vase attributes. No easy way to list only amphoras etc.
B.7	Improved sorting and searching. Thumbnails for every piece. Perhaps they could be on the main page.
B.8	
#	Subject : P08-M-S15
B.1	Used links from the index.
B.2	Checked a reasonable number of vases to see if they satisfy criteria to support or

	debase any of the answers.
B.3	I have a good feeling for what an ancient Greek vase may look like and some idea how large they were.
B.4	Some features which would have helped were disabled or not implemented.
B.5	Pleasant to use, wealth of information. Once implemented, query systems will be very useful.
B.6	Was clear how to make immediate comparisons of vases without opening multiple windows.
B.7	
B.8	
#	Subject : P09-M-S16
B.1	– Clicked on a vase – Waited for info to appear – Looked at picture – Looked at more pictures if necessary – Back to main
B.2	– Basing on previous answers – Description of vases on front page – Assumption & judgment
B.3	– Greeks – Geometric – Based on men, women, animals – Conveyed emotions happiness (wedding) sadness (funeral) – religious – not colorful
B.4	n/a
B.5	– Results screen – Short description shown first – Pictures
B.6	– Slow! (Server) – Not all vases had pictures.
B.7	– Make front page more visual (have an image of vase next to description)
B.8	– ugh! System was ungodly slow
#	Subject : P10-F-S18
B.1	I just want to see the image.
B.2	It’s a long process because the network is too slow.

B.3	Nothing.
B.4	Because I don't want to waste my time.
B.5	I don't like it because it takes me too much time.
B.6	The same as B.5
B.7	I would like to see all the pictures at one time.
B.8	No

D.4.3 PART C : QUIZ

3D Interactive Vase Museum		
Subjects	Score (50)	Skipped
V01-M-S01	27	12
V02-M-S02	44	0
V03-M-S04	30	4
V04-F-S08	37	0
V05-F-S11	39	0
V06-M-S12	37	0
V07-M-S13	39	0
V08-F-S17	36	0
V09-F-S19	35	0
V10-M-S20	37	0
N	10	
Sum	361	
Mean	36.1	
Min	44	
Max	27	
Variance	22.5444	
Stdev	4.7481	
T-test p value	0.3751	

Tufts Perseus Digital Library		
Subjects	Score (50)	Skipped
P01-M-S03	38	0
P02-M-S05	39	3
P03-F-S06	32	0
P04-F-S07	36	0
P05-F-S09	31	0
P06-M-S10	14	34
P07-M-S14	38	0
P08-M-S15	40	0
P09-M-S16	37	0
P10-F-S18	30	0
N	10	
Sum	335	
Mean	33.5	
Min	40	
Max	14	
Variance	59.1667	
Stdev	7.6920	

REFERENCES

[1] C. Ahlberg and B. Shneiderman. (1994a). "**The Alphaslider: A Compact and Rapid Selector**," In Proceedings of CHI'94, ACM Conference on Human Factors in Computing Systems. 365-371.

[2] C. Ahlberg and B. Schneiderman. (1994b). "**Visual Information Seeking Using the FilmFinder**," In ConferenceCompanion of CHI'94, ACM Conference on Human Factors in Computing Systems. 433

[3] C. Ahlberg, E. Wistrand. (1995). "**IVEE: An Information Visualization and Exploration Environment**," In Proceedings of IEEE Information Visualization Symposium '95, 63-73

[4] R. B. Allen (1994). "**Navigating and Searching in Hierarchical Digital Library Catalogs**," In Proceedings of DL'94, ACM Conference on Digital Libraries. 95-100.

[5] The **AMORE** Project © NEC USA, Inc. (1997). http://www.ccrl.com/amore/cats/sports/html/query.html

[6] K. Andrews and H. Heidegger. (1998). "**Information Slices: Visualising and Exploring Large Hierarchies Using Cascading, Semicircular Discs**," In IEEE Information Visualization Symposium Oct., Late Breaking Hot Topics Proceedings, Raleigh Durham, NC, 9-12,

[7] M. D. Apperley, R. Spenser, and C. Gutwin. (2000). "**The Neighborhood Explorer**," Working Paper 00/3, Dept of Computer Science, University of Waikato, New Zealand.

[8] K. W. Authur, K. S. Booth, C. Ware (1993). "**Evaluating 3D Task Performance for Fish Tank Virtual Worlds**," In ACM Transactions on Information Systems, 11(3): 239-265.

[9] R. Baecker and I. Small. (1990). "**Animation at the Interface**," In B. Laurel, editor, The Art of Human-Computer Interface Design, Addison-Wesley, Reading, MA. 251-267.

[10] R. M. Baecker, J. Grudin, W.A.S. Buxton, and S. Greenberg. (1995). "**Readings in Human-Computer Interaction: Toward the Year 2000**," Morgan Kaufmann, San Francisco, ISBN 1-55860-246-1.

[11] L. J. Bannon and S. Bödker. (1991). "**Beyond the Interface: Encountering Artifacts in Use**," In J.M. Carroll (Ed.), Designing Interaction: Psychology at the Human Computer Interface. Cambridge, UK: Cambridge University Press. 227-253.

[12] L. Bartram, A. Ho, J. Dill, and F. Henigman. (1995). "**The Continuous Zoom: A Constrained Fisheye Technique for Viewing and Navigating large Information Spaces**," In Proceedings of ACM UIST '95, 207-215, ACM Press.

[13] R. J. Beach and M. C. Stone. (1983). "**Graphical Style-Towards High Quality Illustrations**," In Proceedings of SIGGRAPH'83, ACM Conference on Computer Graphics. 127-135.

[14] R. A. Beacker and W. S. Cleveland. (1984). "**Brushing a Scatterplot Matrix: High-Interaction Graphical Methods for Analyzing Multidimensional Data**," Technical Report. AT&T Bell Laboratories.

[15] D. Beard and J. Q. Walker II. (1990). "**Navigational Techniques to Improve the Display of Large Two-Dimensional Spaces**," Behavior and Information Technology, 9(6): 451-466.

[16] M. Beard. (1991). "**Looking at Greek Vases**," Adopting an Approach II, In T. Rasmussen & N. Sivey (Eds.), Cambridge University Press, 12-25.

[17] R. A. Becker and W. S. Cleveland. (1987). "**Brushing Scatterplots**," Technometrics, 29(2): 127-142.

[18] R. A. Becker, S. G. Eick, E. E. Miller, and A. R. Wilks. (1989). "**Dynamic Graphics Analysis of Maps**," Holdel, NJ: Bell Laboratories ITV Facility.

[19] B. Bederson and J. Hollan. (1994). "**Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics**," In Proceedings of ACM UIST '94, ACM Press.

[20] B. Bederson, J. Hollan, K. Perlin, J. Meyers, D. Bacon, and G. Furnas. (1996). "**Advances in the Pad++ Zoomable Graphics Widget**," Journal of Visual Languages and Computing, No. 7.

[21] B. B. Bederson. (1994). "**Pad++: Advances in Multiscale Interfaces**," In Proceedings of CHI'94 ACM Conference on Human Factors in Computing Systems, NY. 315-316.

[22] I. Benbasat and P. Todd. (1993). "**An Experimental Investigation of Interface Design Alternatives: Icon vs. Text and Direct Manipulations vs. Menus**," Int'l Journal of Man-Machine Studies. 38:369-402.

[23] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. (1993). "**Toolglass and Magic Lenses: The See-Through Interface**," Computer Graphics (Procs. SIGGRAPH '93), 27(Annual ConferenceSeries):73-80.

[24] S. Björk. (2000). "**Hierarchical Flip Zooming: Enabling Parallel Exploration of Hierarchical Visualizations**," In Proceedings of AVI 2000, ACM Press, 232-237.

[25] S. Björk and L.E. Holmquist. (1998). "**Formative Evaluation of a Focus+Context Visualization Technique**," In Proceedings of HCI '98 (poster), The British HCI Society, Sheffield, UK.

[26] S. Björk and L.E. Holmquist. (1998). "**The Digital Variants Browser: An Explorative Tool for Literature Studies**," In Proceedings of Computers, Literature and Philology, Edinburgh, UK.

[27] S. Björk, L. E. Holmquist, P. Ljungstrand, and J. Redström. (2000). "**Power-View: Structured Access to Integrated Information on Small Screens**," In Extended Abstracts of CHI '00, ACM Press, 265-266.

[28] S. Björk, L. E. Holmquist, and J. Redström. (1999). "**A Framework for Focus+ Context Visualization**." In Proceedings of Information Visualization, © IEEE, 53-56.

[29] S. Björk and J. Redström. (2000). "**Redefining the Focus and Context of Focus+ Context Visualizations**," In Proceedings of IEEE Information Visualization, The Interactive Institute, Sweden.

[30] **Bolles Collection**. (2000). Tufts University Archives, Tufts University, Medford, MA. http://www.library.tufts.edu/archives/london/tallis.htm

[31] R. Bolt. (1984). "**The Human Interface: Where People and Computers Meet**," Wadsworth, Inc.

[32] K. W. Boyack, B. N. Wylie, and G. S. Davidson. (2001a). "**Information Visualization, Human-Computer Interaction, and Cognitive Psychology: Domain Visualizations**," Sandia National Laboratories, Albuquerque, NM. ACM 1-58113-000-0/00/0000.

[33] K. W. Boyack, B. N. Wylie, G. S. Davison. (2001b). "**A Call to Researchers: Digital Libraries Need Collaboration Across Disciplines**," Sandia National Laboratories, D-Lib Magazine Oct.

[34] F. P. Brooks. (1986). "**Walkthrough-A Dynamic Graphics System for Simulating Virtual Buildings**," In Proceedings of the 1986 ACM Workshop on 3D Graphics, Chapel Hill, NC, Oct. 23-24, 9-21.

[35] S. K. Card, G. G. Robertson, and W. York. (1996). "**The WebBook and the WebForager: An Information Workspace for the World-Wide Web**," In Proceedings of CHI'96, ACM, Inc, 111-117.

[36] S. K. Card and J. D. Mackinlay. (1997). "**The Structure of the Information Visualization Design Space**," In Proceedings of InfoVis'97, IEEE Symposium on Information Visualization, Phoenix. 92-99.

[37] M. S. T. Carpendale, D. J. Coperthwaite, and F. D. Fracchia. (1997). "**Extending Distortion Viewing from 2D to 3D**," IEEE Computer Graphics and Applications, Jul., Aug.

[38] M. S. T. Carpendale, D. J. Cowperthwaite, and F. D. Fracchia. (1996). "**Distortion Viewing Techniques for 3D Data**," In Proceedings of InfoVis'96, IEEE Symposium on Information Visualization, Los Alamitos, CA. color plate 119, 46-53.

[39] D. A. Carr, N. Jog, and H.P. Kumar. (1994). "**Using Interaction Object-Graphs to Specify and Develop Graphical Widgets**," Technical Report CS-TR-3344. University of Maryland.

[40] J. Carrière and R. Kazman, (1995). "**Interacting the Huge Hierarchies: Beyond Cone Trees**," In Proceedings of InfoVis'95, IEEE Symposium on Information Visualization, NY. 74-81.

[41] M. Chalmers, R. Ingram, and C. Pfranger. (1996). "**Adding Image-Ability Feature to Information Displays**," In Proceedings of UIST'96, ACM Symposium on User Interface Software and Technology. 33-39.

[42] M. Chalmers. (1993). "**Using a Landscape Metaphor to Represent a Corpus of Documents**," In Proceedings of COSIT '93. Springer, 377-390.

[43] B.-W. Chang and D. Ungar. (1993). "**Animation: From Cartoons to the User Interface**," In Proceedings of UIST'93, ACM Symposium on User Interface Software and Technology, Atlanta. 45-55.

[44] C. Chen. (1998). "**Augmenting User Interfaces for Digital Libraries with Virtual Reality**," In Proceedings 31st Annual Hawaii Int'l Conference on System Sciences, © IEEE,1060-3425/98,

[45] C. Chen. (1999). "**Information Visualisation and Virtual Environments**," Springer.

[46] H.C. Chen and K. T. Chan. (1990). "**Reading Computer-Displayed Moving Text with and without Self-Control Over the Display Rate**," Behavior and Information Technology, 9(6): 467-477.

[47] M. H. Chignell, S. Zuberec, and F. Poblete. (1993). "**An Exploration in the Design Space of Three Dimensional Hierarchies**," In Proceedings of Human Factors Society '93, Santa Monica, CA. 333-337.

[48] R. Chimera and B. Shneiderman. (1994). "**An Exploratory Evaluation of Three Interfaces for Browsing Large Hierarchical Tables of Contents**," ACM Transactions on Information Systems, 12(4 Oct.). 383-406.

[49] M. C. Chuah, S. F. Roth, J. Mattis, and J. "**A. Kolojejckick. (1995). SDM: Selective Dynamic Manipulation of Visualizations**," In Proceedings of UIST'95, ACM, Inc., 61-70.

[50] G. Colby and L. Scholl. (1991). "**Transparency and Blur as Selective Cues for Complex Visual Information**," Image Handling and Reproduction Systems Integration, In SPIE 1460: 114-125.

[51] **Cortona VRML Client**. (2003). Parallel Graphics : http://www.parallelgraphics.com/products/cortona/

[52] G. R. Crane. (Ed.). (1992). "**Perseus 1.0: Sources and Studies on Ancient Greek Culture**," Yale University Press, New Haven. CD-ROM

[53] G. R. Crane. (Ed.). (2000). "**The Perseus Project**. http://www.perseus.tufts.edu, Tufts University, Medford, MA.

[54] G. R. Crane. (1996**). "Building a Digital Library: the Perseus Project as a Case Study in the Humanities**," In Proceedings of 1st ACM Int'l Conference on Digital Libraries, 3-10.

[55] G. Crane and E. Mylonas. (1998). "**The Perseus Project: An Interactive Curriculum on Classical Greek Civilization**," Educational Technology, 28. 25-32.

[56] G. S. Davidson, B. Hendrickson, D. K. Johnson, C. E. Meyers, and B. N. Wylie. (1998). "**Knowledge Mining with VxInsight: Discovery Through Interaction**," Journal of Intelligent Information Systems 11, 259-285.

[57] **Discrete© 3ds max**. (2001). User References, Version 4, Discrete © Autodesk, Inc. www.discrete.com

[58] J. Döllner and K. Hinrichs. (1997). "**Object-Oriented 3D Modeling, Animation and Interaction**," In Journal of Visualization and Computer Animation, 8(1):33-64.

[59] J. Döllner and K. Hinrichs. (2000) "**Dynamic 3D Maps and Their Texture-Based Design**," Institut fűr Informatik, University of Munster, © IEEE 0-7696-0643-7/00.

[60] W. Donelson. (1978). "**Spatial Management of Information**." In Proceedings of SIGGRAPH'78, ACM Conference on Computer Graphics. 203-209.

[61] R. Dristof and A. Satran. (1995). "**Interactivity by Design**," Adobe Press.

[62] K. M. Fairchild, S. E. Poltrock, and G. W. Furnas. (1988). "**SemNet: Three-Dimensional Representations of Large Knowledge Bases**," In R. Guindon (Ed.), Cognitive Science and Its Applications for Human-Computer Interaction. Hillsdale, NJ: Lawrence Erlbaum Associates. 201-233.

[63] S. K. Feiner. (1988). "**Seeing the Forest for the Trees: Hierarchical Display of Hypertext Structures**." In Proceedings of of COIS'88, Palo Alto, CA. 205-212.

[64] S. K. Feiner. (1990). "**Visualizing n-Dimensional Virtual Worlds with n-Vision**," Computer Graphics 24(2), 37-38.

[65] S. K. Feiner and C. Beshers. (1990). "**Worlds within Worlds: Metaphors for Exploring n-Dimensional Virtual Worlds**," In Proceedings of UIST'90, ACM Symposium on User Interface Software and Technology. 76-83.

[66] J. Foley, W.C. Kim, S. Kovacevic, and K. Murray. (1989). "**Defining Interfaces at a High Level of Abstraction**," © IEEE Software, 6(1): 25-32.

[67] J. D. Foley, A. van Dam, S.K. Feiner, and J.F. Hughes. (1996). "**Computer Graphics: Principles and Practice**," Second Edition in C, Addison-Wesley, Reading, MA. ISBN 0 201 84840 6.

[68] G. W. Furnas. (1986). "**Generalized Fisheye Views**," In Proceedings of ACM CHI'86 Conference on Human Factors in Computing Systems, 16-23.

[69] G. W. Furnas. (1999). "**The FISHEYE View: A New Look at Structured Files**," In Readings in Information Visualization: Using Vision to Think, Morgan Kaufmann, San Francisco, CA, 312-330.

[70] G. W. Furnas. (1997). "**Effective View Navigation**," In Proceedings of CHI'97, ACM Conference on Human Factors in Computing Systems, Atlanta. 367-374.

[71] G. W. Furnas and B. B. Bederson. (1995). "**Space-Scale Diagrams: Understanding Multiscale Interfaces**," In Proceedings of CHI '95, ACM Press, 234-241.

[72] K. Garland. (1994). "**Mr. Beck's Underground Map: a History**," Harrow Weald, Capital Transport Publishing.

[73] S. Gauch, R. Aust, J. Evans, J. Gauch, G. Minden, D. Niehaus, and J. Roberts. (1994). "**The Digital Library System: Vision and Design**," In Proceedings of DL'94, ACM Conference on Digital Libraries, 47-52.

[74] N. D. Gershon, S. G. Eick, and S. K. Card. (1998). "**New Visions of Human-Computer Interaction**," ACM Interactions, Mar., 8-15.

[75] V. Giller, M. Tscheligi, J. Schrammel, P. Fröhlich, B. Rabl, R. Kosara, S. Miksch, and H. Hauser. (2001). "**Experimental Evaluation of Semantic Depth of Field, A Preattentive Method for Focus+Context Visualization**," Technical Report TR-VRVis 2001-021, VRVis Research Center, http://www.vrvis.at/, Austria.

[76] M. Green and R. J. K. Jacob. (1991). "**Software Architectures and Metaphors for Non-WIMP User Interfaces**," Computer Graphics, Jul., 25(3): 229-235.

[77] M. Green and S. Halliday. (1996). "**Geometric Modeling and Animation System for Virtual Reality**," Comm. ACM, May, 39(5): 46-53.

[78] D. Gunkel. (2000). "**Rethinking Virtual Reality: Simulation and the Deconstruction of the Image**," Critical Studies in Media Communication, 17(1), 45-62.

[79] P. Haeberli and M. Segal. (1993). "**Texture Mapping as a Fundamental Drawing Primitive**." In Proceedings of the 4th Eurographics Workshop on Rendering, M. Cohen, Puech, C., Sillion, F. (Eds).

[80] V. J. Harward. (1988). "**From Museum to Monitor: The Visual Exploration of the Ancient World**," Academic Computing, 2. 16-19, 69-71

[81] A. J. Head. (2000). "**Design Wise: A Guide for Evaluating the Interface Design of Information Resources**," 2nd Printing, Information Today, Inc.

[82] M. A. Hearst. (1995). "**TileBars: Visualization of Term Distribution Information in Full Text Information Access**," In Proceedings of CHI'95, ACM Conference on Human Factors in Computing Systems. 59-66.

[83] A. Heinonen, S. Pulkkinen, and I. Rakkolainen. (2000). "**An Information Database for VRML Cities**," Ditital Media Institute, Tampere University of Technology, IEEE 0-7695-0743-3/00.

[84] D. A. Henderson, Jr., and S. K. Card. (1986). "**Rooms: The Use of Multiple Virtual Workspaces to Reduce Space Contention in a Window-Based Graphical User Interface**," ACM Transactions on Graphics, 5(3 Jul.), 211-243.

[85] R .J. Hendley, N. S. Drew, A. M. Wood, and R. Beale. (1995). "**Narcissus: Visualizing Information**," In Proceedings of InfoVis'95, IEEE Symposium of Information Visualization, NY. 90-96.

[86] K. P. Herndon, A. van Dam, and M. Gleicher. (1994). "**The Challenges of 3-D Interaction**," In Proceedings of CHI'94, ACM Conference on Human Factors in Computing Systems. 469.

[87] C. F. Herot, R. Carling, M. Friedell, and D. Framlich. (1980). "**A Prototype Spatial Data management System**," Computer Graphics, 14(1), 63-70.

[88] C. F. Herot. (1980). "**Spatial management of Data**," ACM Transactions on Database Systems, 5(4), 493-514.

[89] W. C. Hill and J. D. Hollan. (1994). "**History-Enriched Digital Objects**," The Information Society, 10(2).

[90] E. D. Hoadley. (1990). "**Investigating the Effects of Color**," Communications of the ACM, 33(2), 120-139.

[91] L. E. Holmquist. (1998). "**The Zoom Browser: Showing Simultaneous Detail and Overview in Large Documents**," ITH, Bôras, Sweden, In Human IT. 2(3): 131-150,

[92] L. E. Holmquist and S. Björk. (1998). "**A Hierarchical Focus+Context Method for Image Browsing**," In SIGGRAPH 98 Sketches and Applications, 282, ACM Press.

[93] L. E. Holmquist. (1997). "**Focus+Context Visualization with Flip Zooming and Zoom Browser**." In Proceedings of Conferenceof Human Factors in Computing Systems (Atlanta, Georgia USA). ACM CHI'97, 263-264.

[94] S. E. Hudson. (1989). "**Graphical Specification of Flexible User Interface Displays**." In Proceedings of ACM UIST'89 Symposium on User Interface Software and Technology, Addison-Wesley/ACM Press, Williamsburg, VA, 105-114.

[95] J. E. Hughes. (1998). "**The Perseus Project: Implications for Learning and Teaching**," Ph.D. Educational Psychology, Michigan State University. Able Media.

[96] R. J. K. Jacob, L. Deligiannidis, and S. Morrison. (1999). "**A Software Model and Specification Language for Non-WIMP User Interfaces**," ACM Transactions on Computer-Human Interaction, Mar., 6(1): 1-46.

[97] R. J. K. Jacob. (1993). "**Eye Movement-Based Human-Computer Interaction Techniques: Toward Non-Command Interfaces. in Advances in Human-Computer Interaction**," Ed. by H.R. Hartson and D. Hix, Ablex Publishing Co., Norwood, N.J., 4: 151- 190.

[98] R. J. K. Jacob. (1991). "**The Use of Eye Movements in Human-Computer Interaction Techniques: What You Look At is What You Get**," ACM Transactions on Information Systems, Apr., 9(3): 152-169.

[99] R. J. K. Jacob, H.Y. Shiaw, and G. R. Crane. (2001). "**New Human-computer Interaction Techniques for the Digital Library**," Technical Report 01-1, Jan, Dept. of Computer Science, Tufts University, Medford, MA.

[100] R. J. K. Jacob. (1986). "**A Specification Language for Direct Manipulation User Interfaces**," ACM Transactions on Graphics, 5(4): 283-317.

[101] **Java 3D API.** (2003). Sun Microsystems. http://java.sun.com/products/java-media/3D/

[102] D. F. Jerding and J. T. Stasko. (1995a). "**The Information Mural: A Technique for Displaying and Navigating Large Information Spaces**," In Proceedings of InfoVis'95, IEEE Symposium on Information Visualization, NY. 43-50.

[103] D. F. Jerding and J. T. Stasko. (1995b). "**Using Information Murals in Visualization Applications**," In Proceedings of UIST'95, ACM Symposium on User Interface Software and Technology, Pittsburgh. 73-74.

[104] N. Jog and B. Shneiderman. (1995). "**Information Visualization with Smooth Zooming on a Starfield Display**," In Proceedings of IFIP Conference on Visual Databases, London 1-10.

[105] B. Johnson and B. Shneiderman. (1991). "**Tree-maps: A Space Filling Approach to the Visualization of Hierarchical Information Structures**," In Proceedings of the IEEE Visualization '91 (Oct.), San Diego, CA, 284-291.

[106] C. Johnson. (1998). "**On the Problems of Validating Desktop VR**," In H. Johnson, L. Nigay & C. Roast (Eds.), People and Computers XIII (Procs. of HCI'98). Springer: London. 327-338.

[107] J. A. Johnson, B. A. Nardi, C. L. Zarmer, and J. R. Miller. (1993). "**Information Visualization Using 3D Interactive Animation**," Communications of the ACM, 36(4), 40-56.

[108] T. Keahey. (1998). "**The Generalized Detail-In-Context Problem**," In Proceedings of IEEE Visualization '98, Information Visualization Symposium, IEEE Press.

[109] P. R. Keller and M. M. Keller. (1993). "**Visual Cues: Practical Data Visualization**," Los Alamitos, CA: IEEE Press

[110] C. Keshkin and V. Vogelmann. (1997). "**Effective Visualization of Hierarchical Graphs with the Cityscape Metaphor**," In Proceedings of the Workshop on "New Paradigms in Information Visualization and Manipulation". ACM, 52-57.

[111] H. Kim and R. R. Korfhage. (1994). "**BIRD: Browsing Interface for the Retrieval of Documents**," In Proceedings of the IEEE Symposium on Visual Languages'94. 176-177.

[112] R. Kling and M. Elliott. (1994). "**Digital Library Design for Usability**," Dept of Information and Computer Science Center for Research on Information Technology in Organization, University of California, CA.

[113] T. Kohonen. (1989). "**Self-organization and Associate Memory**," 3rd Ed. Berlin: Springer-Verlag.

[114] T. Kohonen. (1990). "**The Self-Organizing Map**," In Proceedings of the IEEE 78(9), 1464-1480.

[115] H. Koike. (1994). "**Fractal Views: A Fractal-Based Method for Controlling Information Display**," ACM Transactions on Information System, 13(3 Jul.), 305-323.

[116] R. R. Korfhage and K. Olsen. (1991). "**Information Display: Control of Visual Representations**." In Proceedings of IEEE Workshop on Visual Lnaguages'91.

[117] F. Korn. (1995). "**A Taxonomy of Browsing Methods: Approaches to the "Lost in Concept Space" Problem**," Technical Report. University of Maryland, Dept. of Computer Science.

[118] F. Korn and B. Shneiderman. (1995). "**Navigating Terminology Hierarchies to Access a Digital Library of Medical Images**," Technical Report 96-01. University of Maryland, Dept. of Computer Science.

[119] S. M. Kosslyn. (1994). "**Image and Brain: The Resolution of the Imagery Debate**," Cambridge, MA: MIT Press.

[120] L. Koved and B. Shneiderman. (1986). "**Embedded Menus: Selecting Items in Context**," Communications of the ACM, 29(4), 312-318.

[121] M. Kraak. (1994). "**Interactive Modeling Environment for Three-Dimensional Maps: Functionality and Interface Issues**," A. M. MacEachren, D.R., Fraser Taylor (Eds), In Visualization in Modern Cartography, Pergamon, 269-285.

[122] J. Lamping and R. Rao. (1994). "**Laying Out and Visualizing Large Trees Using a Hyperbolic Space**," In Proceedings of UIST'94, ACM Symposium on User Interface Software and Technology. 13-14.

[123] J. Lamping and R. Rao. (1996). "**The Hyperbolic Browser: A Focus+Context Technique for Visualizing Large Hierarchies**," Journal of Visual Languages and Computing, 7(1), 33-55.

[124] J. Lamping, R. Rao, and P. Pirolli. (1995). "**A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies**," Inc Proceedings of CHI '95, ACM Press. 401-408.

[125] J. Larkin and H. A. Simon. (1987). "**Why a Diagram is (Sometimes) Worth Ten Thousand Words**," Cognitive Science, 11(1), 65-99.

[126] J. A. Larson. (1992). "**Interactive Software: Tools for Building Interactive User Interfaces**," Yourdon Press Computing Series, Prentices-Hall, Inc. NJ.

[127] H.-C. Lee. (1990). "**Review of Image-Blur Models in a Photographic System Using Principles of Optics**," Optical Engineering, 29(5): 405-421.

[128] Y. K. Leung and M. D. Apperley. (1994). "**A Review and Taxonomy of Distortion-Orientated Presentation Techniques**," ACM Trans. On Computer-Human Interaction, 1(2): 126-160.

[129] H. Levkowitz and G. T. Herman. (1992). "**Color Scales for Image Data**," IEEE Computer Graphics and Applications, 12, 78-80.

[130] H. Levkowitz, R. A. Holub, G. W. Meyer, and A. K. Robertson. (1992). "**Color vs Black and White in Visualization**," IEEE Computer Graphics and Applications, 12(14), 20-22.

[131] X. Lin. (1996). "**Graphical Table of Contents**," In Proceedings of DL'96, ACM Conference on Digital Libraries. 45-53.

[132] X. Lin. (1992). "**Visualization for the Document Space**." In Proceedings of IEEE Visualization '92, © IEEE, 274-281.

[133] G. Lindgaard. (1994). "**Usability Testing and System Evaluation: A Guide for Designing Useful Computer Systems**," Chapman & Hall

[134] A. Lockwood. (1969). "**Diagrams: A Visual Survey of Graphs, Maps, Charts and Diagrams for the Graphic Designer**," NY: Watson-Guptill.

[135] L. W. MacDonald. (1990). "**Using Colour Effectively in Displays for Computer-Human Interface. Displays**," 129-142.

[136] J. D. Mackinlay, G. G. Robertson, S. K. Card. (1991). "**Perspective Wall: Detail and Context Smoothly Integrated**," ACM, Proceedings CHI '91, 173-179.

[137] G. Mar.ionini and G. Crane (1994). "**Evaluating Hypermedia and Learning: Methods and Results from the Perseus Project**," ACM Transaction on Information Systems, (12): 5-35.

[138] A. Marcus. (1991). "**Graphic Design for Electronic Documents and User Interfaces**," NY: ACM Press.

[139] M. McCarthy and A. Descartes. (1988). "**Reality Architecture, Building 3D worlds with Java and VRML**," Prentice Hall Europe.

[140] D. Modjeska and J. Waterworth. (2000). "**Effects of Desktop 3D World Design on User Navigation and Search Performance**," Dept. of Computer Science, University of Toronto, © IEEE 0-7695-0743-3 /00.

[141] A. Morse. (1979). "**Some Principles for the Effective Display of Data,"** In Proceedings of SIGGRAPH'79, ACM Conference on Computer Graphics and Interaction, 13(2), 94-101.

[142] E. Mylonas (1993). "**The Perseus Project**," In Scholarly Publishing on the Electronic Networks: The New Generation: Visions and Opportunities in Not-for-Profit Publishing: Proceedings of the Second Symposium, ed. Ann Okerson, 81-88. Washington, DC: Office of Scientific and Academic Publishing, Association of Research Librarie.

[143] D. Neuman. (1991). "**Evaluating Evolution: Naturalistic Inquiry and the Perseus Project**," Computers and the Humanities, 25. 239-246.

[144] J. Nielsen. (1993). "**Noncommand User Interfaces**," Comm. ACM, 36(4): 83-99.

[145] D. A. Norman. (1988). "**The Psychology of Everyday Things**," Basic Books, NY.

[146] D. A. Norman. (1991). "**Cognitive Artifacts**," In J. Carroll (Ed.), Designing Interaction, Cambridge University Press, NY.

[147] C. North and B.Shneiderman. (1999) "**Snap-Together Visualization: Evaluating Coordination Usage and Construction**", University of Maryland Computer Science Dept. Technical Report #CS-TR-4075.

[148] C. North, B. Shneiderman, and C. Plaisant. (1996). "**User Controlled Overviews of an Image Library: A Case Study of the Visible Human**." In Proceedings of 1st ACM Int'l Conference on Digital Libraries, Mar., 74-82.

[149] K. A. Olsen, R. R. Korhage, K. M. Sochats, M. B. Spring, and J. Williams. (1993). "**Visualization of a Document Collection: The VIBE System**," Information Processing and Management, 29(1), 69-81.

[150] M. Osada, H. Liao, and B. Shneiderman. (1993). "**Alphaslider: Searching Textual Lists with Sliders**," In Proceedings of Ninth Annual Japanese Conference on Human Interface '93.

[151] K. Perlin and D. Fox. (1993). "**Pad: An Alternative Approach to the Computer Interface**," In Proceedings of SIGGRAPH '93, ACM Press.

[152] K. Pimentel and K. Teixeira. (1993). "**Virtual Reality: Through the New Looking Glass**," 1st Edition, Intel® / Windcrest® / McGraw-Hill Inc.

[153] P. Pirolli and R. Roa (1996). "**Table Lens as a Tool for Making Sense of Data**." In Proceedings of ACM CHI'96 (Vancouver, Apr. 1996). ACM Press, 67-80.

[154] C. Plaisant, D. Carr, and B. Shneiderman. (1995). "**Image-Browser Taxonomy and Guidelines for Designers**," IEEE Software, 12(2):21-32.

[155] C. Plaisant, A. Rose, B. Milash, S. Widoff, and B. Shneiderman. (1996). "**LifeLines: Visualizing Personal Histories**." In Proceedings of ACM CHI96 Conf., Apr., 221-227, 518.

[156] S. Pook, E. Lecolinet, G. Vaysseix, E. Barillot. (2000). "**Context and Interaction in Zoomable User Interfaces**," In Proceedings of AVI, ACM Press, 227-231.

[157] J. Preece, D. Benyon, G. Davies, L. Keller, and Y. Rogers. (1993). "**A Guide to Usability: Human Factors in Computing**," Addison-Wesley.

[158] J. Preece, Y. Rogers, H. Sharp. (2002). "**Interaction Design: Beyond Human-Computer Interaction**," John Wiley & Sons, Inc.

[159] B. Preimj, A. Raab, and T. Strothotte. (1997). "**Coherent Zooming of Illustrations with 3D-Graphics and Text**." In Proceedings of Graphics Interface '97m 195-113.

[160] C. Proal, J. A. Sánchez, and L. Fernández. (2000). "**UVA: 3D Representations for Visualizing Digital Collections**," In Proceedings of the 3rd Int'l Conference on Visual Computing (Visual 2000, Mexico City)

[161] R. Rao, J. O. Pedersen, M.A. Hearst, J. D. Mackinlay, S. K. Card, L. Masinter, P.-K. Halvorsen, and G. G. Robertson. (1995**)**. "**Rich Interaction in the Digital Library**," In Communications of the ACM, ACM Press, 38(4): 29-39.

[162] R. Rao and S. K. Card. (1994). "**The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus+Context Visualization for Tabular Information**." In Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems, Information Visualization 1, 318-322, Color plates on 481-482.

[163] "**Readings in Human-Computer Interaction: Toward the Year 2000**," (1995, 2000). Written and edited by: Baecker, Ronald, M. Baecker; Grudin, Jonathan; Buxton, William A.S.; Greenberg, Saul. Morgan Kaufmann Publishers, Inc.

[164] "**Readings in Information Visualization: Using Vision to Think**," (1999). Written and edited by: Stuart K. Card, Jock D. MacKinlay and Ben Shneiderman, Academic Press.

[165] C. Reed-Ballreich. (1996). "**Texture Mapping in VRML 2.0**," Name3D, 8.10.

[166] J. W. Regain, W. L. Shebilske, and J. M. Monk, (1992). "**Virtual Reality: An Instructional Medium for Visual-Spatial Tasks**," Journal of Communication, 42(4), 136-149.

[167] M. Rettig. (1992). "**Interface Design When You Don't Know How**," Communications of the ACM 35, 29-34.

[168] G. G. Robertson, J. D. Mackinlay, S. K. Card. (1995). "**Cone Trees: Animated 3D Visualisations of Hierarchical Information**," ACM, Proceedings of CHI '95, 189-194.

[169] G. G. Robertson, S. K. Card, J. D. Mackinlay. (1993). "**Information Visualization Using 3D Interactive Animation**," Communications of the ACM, 36(4): 57-71.

[170] G. G. Robertson and J. D. Mackinlay. (1997). "**The Document Lens**," In Proceedings of UIST'93, © IEEE, 101-108.

[171] S. F. Roth and J. Mattis. (1990). "**Data Characterization for Intelligent Graphic Presentation**," In Proceedings of CHI'90, ACM Conference on Human Factors in Computing Systems, NY. 193-200.

[172] S. F. Roth, P. Lucas, J. A. Senm, C. C. Gomberg, M. B. Burks, P. J. Stroffolino, and J.A. Kolojejchick. (1996). "**Visage: A User Interface Environment for Exploring Information**," In Proceedings of InfoVis'96, IEEE Symposium on Information Visualization. 3-10.

[173] Roth, S.F. (1994). The **SAGE** Project, www.cs.cmu.edu/Web/Groups/sage/sage.html/

[174] J. Rountree, W. Wong, and R. Hannah. (2002). "**Learning to Look: Real and Virtual Artifacts**," ISSN 1436-4522. Educational Technology & Society 5(1).

[175] M. Sarkar and M. H. Brown. (1992). "**Graphical Fisheye Views of Graphs**," In Proceedings of CHI'92, ACM Conference on Human Factors in Computing Systems, NY. 83-91.

[176] M. Sarkar and M. H. Brown. (1994) "**Graphical Fisheye Views**," In Communications of the ACM, ACM Press, 37(12): 73-84.

[177] M. Sarkar, S. S. Snibbe, O. J. Tversky, and S. P. Reiss. "**Stretching the Rubber Sheet: A Metaphor for Visualizing Large layouts on Small Screens**," In Proceedings of the ACM Symposium on User Interface Software and Technology, Visualizing Information, 81-91, 1993.

[178] M. Scaife and Y. Rogers. (1996). "**External Cognition: How Do Graphical Representations Work?**" Int'l Journal of Human-Computer Studies, 45(2), 185-213.

[179] B. Shneiderman. (1983). "**Direct Manipulation: A Step Beyond Programming Languages**," IEEE Computer, 16(8), 57-58.

[180] B. Schneiderman. (1992a). "**Tree Visualization with Treemaps: A 2D Space-Filling Approach**," In ACM Transactions on Graphics. 11(1): 92-99.

[181] B. Shneiderman. (1992b). "**Designing the User Interface: Strategies for Effective Human-Computer Interaction**," 2nd Edition, Addison-Wesley Publ. Co., Reading MA.

[182] B. Shneiderman. (1996). "**The Eyes Have it: A Task by Data Type Taxonomy for Information Visualization**," In Proceedings of 1996 IEEE Symposium on Visual Languages, 336-343, Boulder, Colorado, IEEE Computer Society.

[183] B. Shneiderman.(1997). "**Designing the User Interface: Strategies for Effective Human-Computer Interaction**," Third Edition, Addison-Wesley, Reading, MA.

[184] B. Shneiderman and H. Kang. (2000). "**Direct Annotation: A Drag-and-Drop Strategy for Labeling Photos**," Dept of CS, HCI Lab, Institute for Advanced Computer Studies & Institute for Systems Research, University of Maryland.

[185] B. Shneiderman and K. Norman. (1992). "**Questionnaire for User Interface Satisfaction (QUIS), Designing the User Interface: Strategies for Effective Human-Computer Interaction**," 2nd Edition, Addison-Wesley Publ. Co., Reading MA.

[186] H. Small. (1997). "**Update on Science Mapping: Creating Large Document Spaces**," Scientometrics 38, 275-293.

[187] P.A. Smith and J.R. Wilson. (1993). "**Navigating in Hypertext through Virtual Environments**," Applied Ergonomics, 24(4), 271-278.

[188] R. Spence. (1993). "**A Taxonomy of Graphical Presentation**," In INTERACT '93 and CHI'93 ConferenceCompanion, ACM Press, 113-114, 1993

[189] R. Spence. (2001). "**Information Visualization**," Addison-Wesley.

[190] L. Staples. (1993). "**Representation in Virtual Space: Visual Convention in the Graphical User Interface**," In Proceedings of INTER-CHI'93, ACM Conference on Human Factors in Computing Systems. 348-354.

[191] M. C. Stone, K. Fishkin, and E. A. Bier. (1994). "**The Movable Filter as a User Interface Tool**," In Proceedings of CHI'94, ACM Conference on Human Factors in Computing Systems, NY. 306-312.

[192] I. E. Sutherland. (1965). "**The Ultimate Display**." In Proceedings of IFIP'65, Conferenceof Int'l Federation for Information Processing. 506-508

[193] V. Tanriverdi and R. J. K. Jacob. (2000). "**Interacting with Eye Movements in Virtual Environments**," In Proceedings of ACM CHI'00 Human Factors in Computing Systems Conf., Addison-Wesley/ACM Press, 265-272.

[194] A. Treisman. (1985). "**Preattentive Processing in Vision**," Computer Vision, Graphics, and Image Processing, 31:156-1777.

[195] E. R. Tufte. (1990). "**Envisioning Information**," Cheshire, CT: Graphics Press.

[196] J. R. Vacca. (1996). "**VRML: Bringing Virtual Reality to the Internet**," Academic Press.

[197] S. Vallance and P. Calder. (2001). "**Context in 3D Planar Navigation**," School of Informatics & Engineering. The Flinders University of South Australia. © IEEE 0-7695-0969-X/01.

[198] A. Veerasamy, S.E. Hudson, and S. Navathe. (1995). "**Querying, Navigating and Visualizing an Online Library Catalog**," In Proceedings of DL'95, ACM Conference on Digital Libraries.

[199] J. Vince. (1998). "**Essential Virtual Reality Fast**," Springer.

[200] **Virgilio** Project : Visual Metaphors for Database Exploration. (2003) OASYS - Open Adaptive Information Management Systems. http://www.ipsi.fhg.de/oasys/projects/virgilio/

[201] Virtual Reality Modeling Language (**VRML**). 2.0, IS, 4.8 (1996). http://www.vrml.org/Specifications/VRML 2.0

[202] **VRoom** Project. MITRE Corporation © (2002). http://www.mitre.org/pubs/showcase/2002/vroom/vroom-prn.html

[203] C. Ware, K. Arthur and K.S. Booth. (1993). "**Fish Tank Virtual Reality**," In Proceedings of ACM INTERCHI'93 Human Factors in Computing Systems Conf., Addison-Wesley/ACM Press, 37-41.

[204] C. Ware and M. Lewis. (1995). "**The DragMag Image Magnifier**," In Video Proceedings of ACM CHI'95, 2: 407-408.

[205] C. Ware. (2000). "**Information Visualization: Perception for Design**," Academic Press.

[206] C. Williamson and B. Shneiderman. (1992). "**The Dynamic HomeFinder: Evaluating Dynamic Queries in a Real-Estate Information Exploration System**." In Proceedings of ACM SIGIR'92 Conf., Copenhagen, Denmark, 338-346.

[207] N. Yankelovich, N. Meyrowitz, and A. van Dam. (1985). "**Reading and Writing the Electronic Book**," IEEE Computer 18(10): 15-30.