The 3D Vase Museum:
A New Approach to Context in a Digital Library

Horn-yeu Shiaw
Tufts University
Computer Science Department
Medford, Mass, U.S.A.
1-617-627-3217

hshiaw@cs.tufts.edu

Robert J.K. Jacob
Tufts University
Computer Science Department
Medford, Mass, U.S.A.
1-617-627-3217

jacob@cs.tufts.edu

Gregory R. Crane
Tufts University
Classics Department
Medford, Mass, U.S.A.
1-617-627-3830

gcrane@perseus.tufts.edu

ABSTRACT

We present a new approach to displaying and browsing a digital library collection, a set of Greek vases in the Perseus digital library. Our design takes advantage of three-dimensional graphics to preserve context even while the user focuses in on a single item. In a typical digital library user interface, a user can either get an overview for context or else see a single selected item, sacrificing the context view. In our 3D Vase Museum, the user can navigate seamlessly from a high level scatterplot-like plan view to a perspective overview of a subset of the collection, to a view of an individual item, to retrieval of data associated with that item, all within the same virtual room and without any mode change or special command. We present this as an example of a solution to the problem of focus-plus-context in information visualization. We developed 3D models from the 2D photographs in the collection and placed them in our 3D virtual room. We evaluated our approach by comparing it to the conventional interface in Perseus using tasks drawn from archaeology courses and found a clear improvement. Subjects who used our 3D Vase Museum performed the tasks 33% better and did so nearly three times faster.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation (e.g., HCI)**]: User Interfaces – *interaction styles, prototyping, user-centered design, evaluation/methodology.*

General Terms

Performance, Design, Experimentation, Human Factors.

Keywords

Human-computer Interaction (HCI), Information Visualization, Virtual Reality (VR), Perseus Digital Library, Usability, Focus versus Context, Focus-plus-Context (F+C).

1. INTRODUCTION

The Perseus digital library [9] contains, among many collections, a set of photographs of Greek vases with accompanying text data about each one. Using a browser, you can query the system and receive a list of links or thumbnails for any selected subset of the vases. You can then visit each link and view the photograph(s) and accompanying information for that vase. But each time you do so, you lose your context; you can see only one vase at a time. It is difficult to get an overview of the entire collection and see each individual vase within that overview. By contrast, if you visited a physical museum, you could scan a roomful of items and compare and contrast them in the context of the collection. Even when you walk right up to an individual item, you retain some implicit knowledge of where it fits in the larger space (near one end, near the middle), merely by your sense of your location within the room, from proprioceptive feedback, from your memory of how you reached this place, and from the view of the rest of the room in your peripheral vision. If the room were arranged by, for example, date or location, this would give you useful information about the context of the vase you are viewing. By contrast, a conventional web browser display is more like sitting in an empty room at the museum and asking the curator to show you one vase at a time to look at. If you request a new vase, the previous one is removed from view.

In a real museum (and in our virtual museum), you can view the whole collection and navigate through it by walking (or by virtual navigation). In a virtual museum, you could also request that the collection instantly be rearranged for you by date, color, size, or any other parameter you like, and you can wander through it in each such dimension.

In an experiment, we used the vase collection to test our claim that presenting information in a 3D environment can provide better context to the user. We took the Greek vase data in the Perseus Digital Library and built a prototype "3D Vase Museum" virtual room in the form of VRML [22] dataset that the user can view and navigate on the screen in a web browser. We evaluated our approach with a set of tasks drawn from archaeology course assignments and found a distinct improvement in speed and performance compared to a conventional browser interface.

The problem we are solving is an example of the more general problem of focus plus context in information visualization. There is a constant trade-off between '*a view of the whole data available, while pursuing detailed analysis of a part of it*' [6]. A standard web interface such as that of Perseus allows learning by

browsing and searching the text and images of the vase collection in Perseus and viewing individual items one by one. It epitomizes the issue of focus versus context in a digital library. Our approach instead allows a user to focus on individual items without losing his or her overall context in the collection.

Our 3D Vase Museum is a virtual geometric landscape arranged in a scatterplot layout in a room-like setting. The data points are 3D models extrapolated from 156 Greek vases in the London collection of the Perseus Digital Library [9]. This dataset was selected for its beauty, controllable data size and non-textual content. Using a natural navigation metaphor, the user can learn by moving within a virtual information space and interacting with a vase (focus) in relation to the London collection (context) in the surroundings.

We found that our approach provides measurable and substantial benefits to the user. To measure this, we compared the 3D Vase Museum to its source Perseus sites in an experiment with two groups of subjects performing tasks based on archaeology course assignments. The results of our experiment showed the group who used the 3D Vase Museum learned nearly three times faster (13.69 vs. 37.03 minutes) with 33% greater accuracy in answering ten task-related questions (7.2 vs. 5.4 correct answers). They also retained information about the collection equally well, despite the fact that they spent only a third as much time learning it. In this paper, we describe the design and implementation of our 3D Vase Museum and present our evaluation experiment.

2. RELATED WORK

Our research builds upon the highly successful Perseus Digital Library Project [9] at Tufts University. Perseus contains many collections, particularly in Greek and Roman literature, of which the London vase collection we used is one. It is viewed with a web browser and a standard graphical user interface (GUI) style of interaction [20]. In developing our new approach to presenting this collection, we draw on work from new 2D and 3D interaction techniques, such as direct manipulation, selective manipulation, zoomable user interfaces, dynamic exploration and interactive exploration, as well as techniques for information visualization [6], especially for focus-plus-context visualization [1, 2, 3, 11, 14, 15, 18]. Our work presents a context-rich information landscape, populated with 3D ancient data artifacts [13], where the user can learn in virtual reality [7] using non-WIMP [12] and lightweight interactions. Our new techniques can be alternatives or additions to the current approaches. What sets our work apart from existing focus-plus-context research is the method by which we were able to validate our approach in usability [16, 19, 21].

3. THE VASE MUSEUM SYSTEM

Figure 1 shows the London collection in the existing Perseus digital library with the 156 Greek vases in a list view. Figure 2 shows the London collection in a thumbnail view, which shows the pictures of the vases in various levels of image quality. A user can select a vase page by clicking on its catalogue number in the list and will receive an HTML page about that vase. Figure 3 shows a sample page for the Greek vase with a catalog number of London 1971.11-1.1.
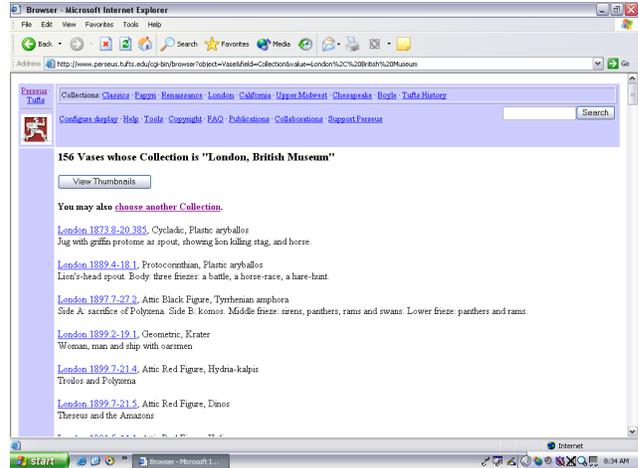
Figure 1. The London Collection in list view.
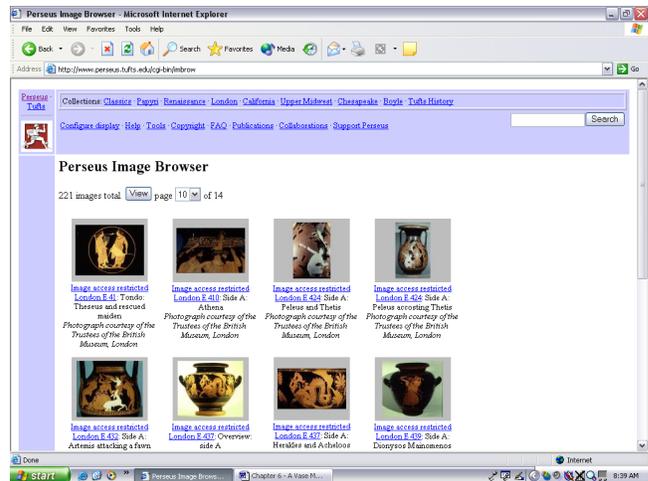
Figure 2. The London collection in thumbnail view.
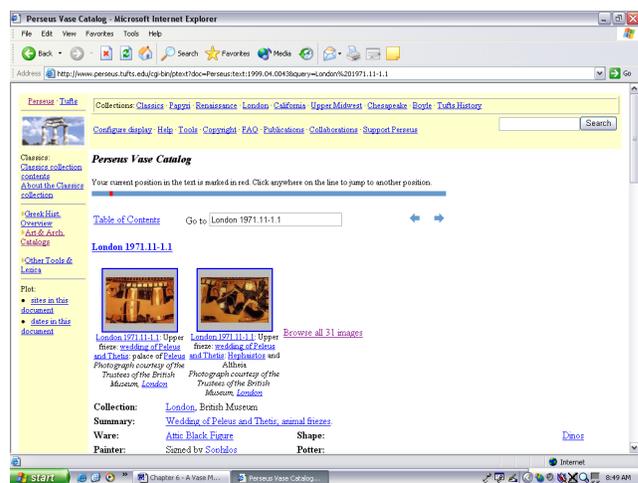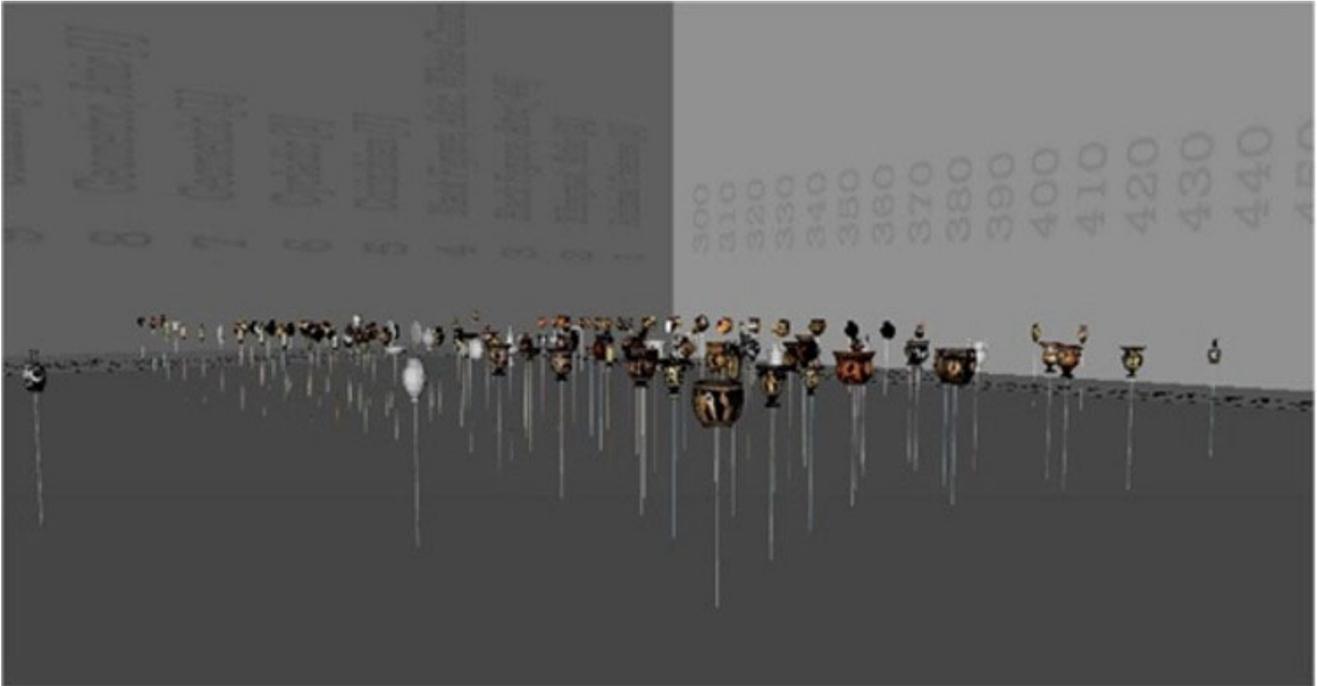
Figure 3. Sample vase London 1971.11-1.1 in page view.

Figure 4. The 3D Vase Museum in wide-angle view.

Our approach replaces the Perseus browser interface with the 3D Vase Museum. The museum is a 3D virtual environment that presents each vase as a 3D object and places it in a 3D space reminiscent of a room within a museum. The user can view the 3D world and navigate anywhere in it with a VRML-enabled web browser or a head-mounted virtual reality display. Figure 4 shows a wide-angle view of the 3D Vase Museum. One wall represents the date (year B.C.E.) and the adjacent wall represents the wares (types of glazing, e.g., red figures, black figures). Figure 5 shows an eye-level view of the 3D Vase Museum. This is the 'normal view' for a user who is walking around in the virtual room and looking at the vases.

Figure 5. The 3D Vase Museum in eye-level view.

In addition to moving around the 3D room and viewing the graphics, we provide a way to integrate other textual information into the view without leaving the context of the 3D room. As the user navigates in X- and Y-axes toward an area of visual interest, secondary information about the vases in the area would appear in the virtual scene. This effect is shown in Figure 6. As the user approached vase London B 376, a text caption appeared on the left and a HTML page appeared on the right. If the user looked closer, the readings became clearer. As he or she moved farther and farther away, the information became less and less visible until it eventually disappeared from the scene.

Figure 6. Secondary information on vase London B 376.

As the user navigated in the vase museum, secondary information appeared and disappeared in the scene triggered by the proximity to the corresponding vases. An animated tour was also available to guide the user in a pre-scripted path. There was a pre-defined camera view for each vase in the scene. A user could select a pre-defined vase view and get a full frontal view of the selected vase (similar to Figure 6). Camera views allowed the user to move

directly to another vase, bypassing virtual walking. If the user clicked on the vase HTML page, a version of the original Perseus HTML page would be loaded, from which a 3D model of the vase could be loaded, as shown in Figure 7. The user could then rotate the individual 3D vase on the screen using the mouse (or an equivalent devise, such as the touchpad on a laptop PC). The 3D Vase Museum is maintained in the background at all times.
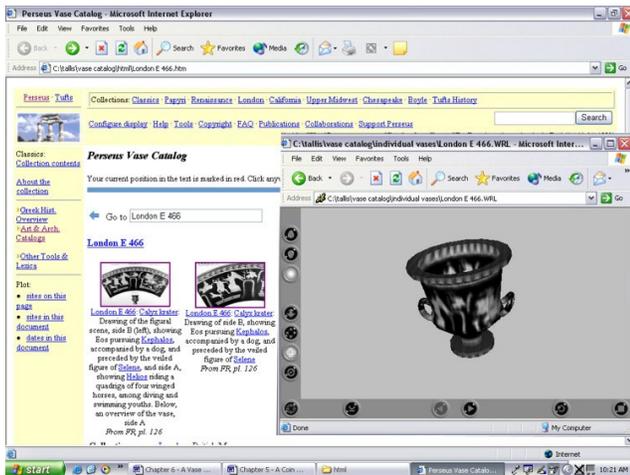
Figure 7. Vase London E 466 launched from a HTML page.

If the user switched the view port upward toward the ceiling (in the Z direction), they would see the room in perspective, as shown in Figure 8. Because the room was arranged by year along one wall and ware along the other, if they looked straight down at the floor from the ceiling, they would see a 2D scatterplot of the vases with years and wares as the axes, as shown in Figure 9. The result is a high-level scatterplot view that you can walk right into without losing your point of view or context — a seamless blend from a high-level data plot (Figure 9) to 3D objects (Figure 6) simply by moving within the virtual environment, without any explicit mode change command or change in frame of reference.
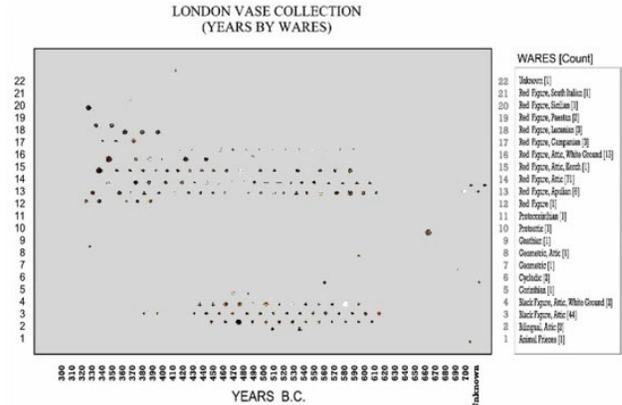
Figure 9. The 3D Vase Museum in 2D scatterplot view.

3.1 Implementation

The 3D Vase Museum was prototyped in components. The 3D vases and various renditions of the museum room were modeled in Discrete© 3ds max 4.3 (*.max) [10], exported to VRML97 [22] file format (*.WRL) and run on the Cortona VRML browser plug-in [8]. Interfaces for the optional immersive VR equipment, such as a head-mounted display (HMD), data gloves and tracking devices, were written in either C, Java or Java3D [17] and called

Figure 8. The 3D Vase Museum in perspective view.

from the VRML scenes as functions. For the experiment, we used the desktop VR version with no immersive VR equipment. The computer was a Toshiba® Satellite 5005-S507 Intel® Pentium® III Laptop PC with 1.10 GHz, NVIDIA® GeForce4 graphic card and 32 MB of VRAM.

The 3D Vase Museum pushed the limit in available PC resources. Much of the prototype implementation reflects a delicate trade-off between virtual realism and rendering speed. The 3D Vase Museum contains 2,155 objects with 260,411 vertices and 510,614 faces. The physical memory usage was between 464.2 MB (minimum at frame 0) and 511.5 MB (maximum). The virtual memory usage was between 417.2 MB (minimum at frame 0) to 1,506.3 MB (maximum). Users were fixed in a walking mode, with a speed of 2.0 and a height of 75 units per second. Scenes in the 3D Vase Museum were optimized for export to VRML97. Animations other than position, rotation and scaling were avoided. Lights were omni. Cameras were free or targeted.

For best rendering performance, a scene should have contained a maximum of 5,000–10,000 polygons with minimal use of texture maps [10], but we pushed beyond that, because each vase was a mesh with a unique image texture map. In an effort to reduce the number of texture maps in the scene, only the best image or side A was used in the modeling if more than one image was available. Primitives, rather than complex meshes, were used whenever possible. The rendering frame rate specified the processor usage in the range from 0 (idle and minimum) to 100 (maximum). The 3D Vase Museum was run at the maximum frame rate. In real time, navigation could appear to be "chunky" when frames were being dropped as the real-time rendering capabilities of the browser were strained. The resulting system was usable but clearly pushed the performance of a current PC to its limits. Throughout the project, we believed that soon after we finished our work, PC technology would catch up to our system, and this appears to be happening.

3.2 Modeling the Room
As shown in Figures 4, 8 and 9, the museum room was relatively large, containing a sparse scatterplot. Only areas with exhibits were lit to discourage the users from navigating to a pitch-black area with no vases. The room itself was an infinite boxed space with an open top in which the user could explore on the floor without ever leaving the room. The floor of the museum was a

floating plane, texture-mapped with the coordinate system. The 3D vases served as the data points in the virtual scatterplot graph. The walls were texture-mapped with the labels of the years and wares. One problem encountered during implementation was the perspective distortion of the walls in alignment of the vases. For example, vases in 500 B.C. lined up against the wall appeared to belong in a different period if viewed from the far left or far right.

Each vase was modeled in a separate file and X-referenced [10] by the various renditions of the 3D Vase Museum. One rendition is shown in Figure 10 below, a screen captured in 3ds max. In this rendition, the 3D vases are lined up in a sequence by their catalog numbers. As shown in Figure 10, the vases were placed on columns, which were thin, semi-transparent cylinders to reducing obscuring of the items behind them. The captions (shown to the left side of the vases) were 3D text (splines). The HTML pages (shown to the right side of the vases) were modeled as very thin, semi-transparent, texture-mapped boxes. The vase captions and HTML pages were billboards (the squares above the vases), which rotate around the corresponding vases and adjust to the direction of the user. Every vase in the museum had a virtual proximity sensor (the diamonds inside the vases) that could be triggered when approached. The secondary information space displayed the caption and HTML page as level of detail (LOD) in proximity. Anchors in the HTML pages loaded the corresponding 2D HTML pages like the one in Perseus shown in Figure 7.

The scene graph of the 3D Vase Museum consisted of the parent room and the children vases X-referenced in the scene. The vases were broken down into the actual vase meshes, plus all supporting objects such as the captions, HTML pages, billboards, lights, cameras, etc. Camera views were preset view ports.

3.3 Modeling the Vases
The London collection contained "vases" were really containers, pitchers, bowls, cups, plates, lids or even an ensemble. The 3D vases were modeled as a form of information art attempting to be true to the nature of the photographic images in Perseus. A 3D object was created from each 2D photograph by assuming the vase generally had the same profile all the way around, as if it were made on a potter's wheel or a lathe. Additional photographs helped model features such as handles or spouts. Many vases had one or two handles. Simple handles were made from bending toruses or cylinders. Elaborate handles were made from meshes.
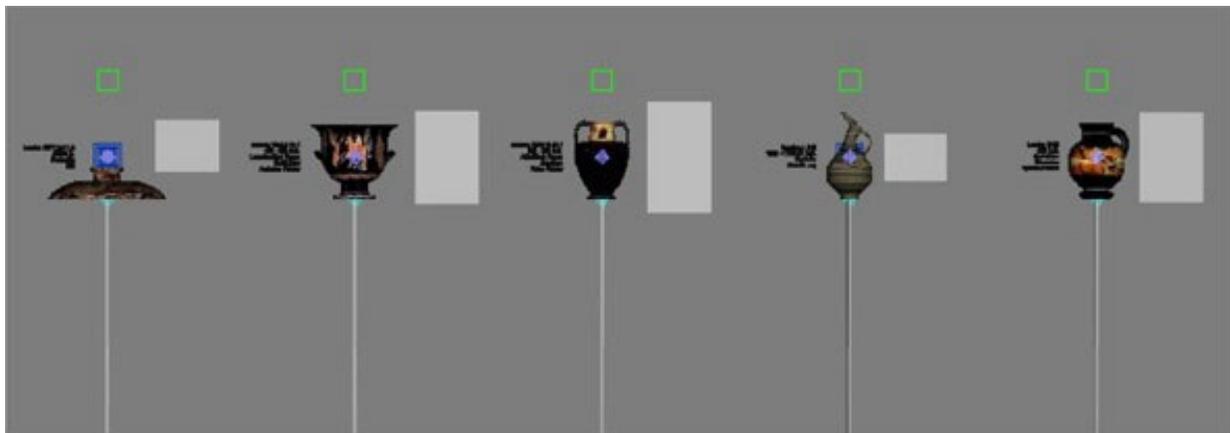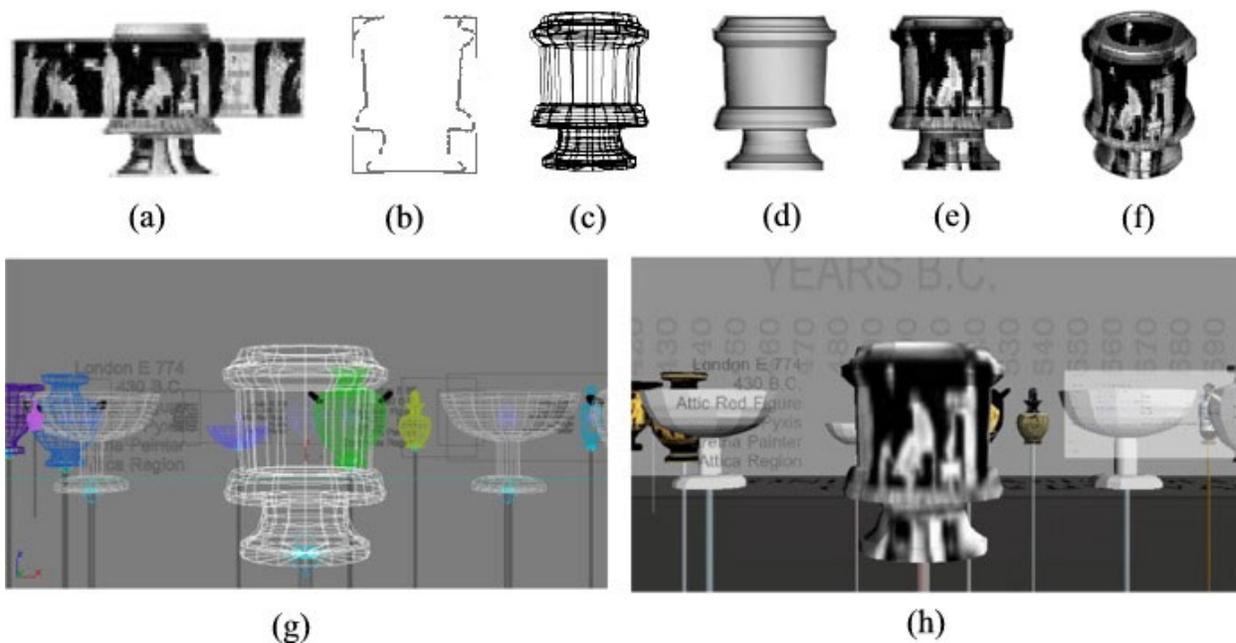
Figure 10. The 3D vases and supporting objects.

Figure 11. Modeling 3D vase London E 774.

The result is a texture-mapped 3D object which the user can then view from any angle and move and rotate at will. Figure 11 shows the process of modeling a sample Greek vase, i.e. London E 774. Given an image of the vase (a) from Perseus, we drew a 2D line contour (b) around the perimeter of (a), built a wire-frame model (c), and shrink-wrapped the lathe model (d) with a JPEG (*.JPG) texture map, created from the original image (a), to get the final product in (e). The perspective view of the 3D model is shown in (f). London E 774 can be seen in the museum setting in wire-frame mode (g) and in 3D perspective mode (h). The 3D vases in the museum sometimes appear to be pixellated when viewed close up, due to the limited resolution of the original images.

Vases with no images were shown as plain grey lathes, similar to the cups next to our sample London E 774 shown in Figure 11(d). In our experiment, the subjects were able to determine on their own, without ever being told the significance of the grey models, the approximate percentage of vases in the collection that had no images (task question #5 below).

4. EVALUATION TASK

The next step is to evaluate the 3D Vase Museum in a user study. A key benefit we claim for our system is its ability to provide overview information about the London collection even while focusing on individual vase items. To measure this, we want to find a task that requires studying the whole collection more than individual vases. The underlying goal of our task is to learn about the 156 vases in the London collection as well and as quickly as possible. Our task should simulate open-ended browsing and studying in a digital library, but in a bounded way within a single experimental session rather than, say, a semester college course.

In order to devise a realistic user task, we examined homework assignments from a variety of college courses in archaeology and developed ten questions modeled after them. These questions reflected the overall context of the London collection, such as the color, themes, shapes, artwork and general appearance of the 156 vases. We simplified the questions to be suitable for a user with no previous study of archaeology, relying only on our digital library. The ten multiple-choice task questions are shown below (correct answers in bold). The answers to the questions were readily available using either conventional Perseus or our system.

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.

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

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.

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

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

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

7. Which color will you NOT find in this collection?
   - **a.) Blue**
   - b.) Black
   - c.) Gold
   - d.) Brown
   - e.) Cannot be determined

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

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.

10. What are the two most common shapes in this collection?
    - **a.) Amphora and Stamnos**
    - b.) Lebes and Loutroforus
    - c.) Kantharos and Volute Krater
    - d.) None of the above
    - e.) Cannot be determined

The subjects were asked to focus on a particular vase or group of vases and compare it with the other vases in the collection. They were asked to observe similarities or differences among the vases. They were told that no reading or counting was required. They were not required to have any prior knowledge of the vases. For task question #10 above, line drawings for the shape terms in the question were shown to the subjects to eliminate the need to know the archaeological terms.

5. EXPERIMENT

We evaluated our 3D Vase Museum by comparing it to the original Perseus web site. Our plan was to design a set of data collection instruments that was equivalent and unbiased to the system in use. Our goal was to evaluate speed and accuracy in learning the general context of the London collection.

5.1 Design

The experiment was conducted in the HCI lab at Tufts. We used a between-subjects design because the task uses general knowledge about the vase collection as a whole, and subjects completing one condition would have learned too much about the collection to be useful in a second condition. We measured subjects' speed and accuracy in answering the ten task questions. For each task question, there were also eight survey questions about how the subject performed that task. After the subjects had completed the ten task questions and associated surveys, they were asked to complete a questionnaire in three parts. Part A was a survey of twenty questions about the experiment. Part B was a section for written comments. Part C was a final quiz with fifty true/false questions similar to the original task questions. It was designed to measure retention and was placed at the end to maximize the interval between performing the task and taking the quiz.

5.1.1 Subjects

The 20 subjects were divided randomly into two groups; the P group used the conventional Perseus system, and the V group used the 3D Vase Museum. Each group consisted of 10 subjects, 6 male and 4 female, generally in their twenties. The subjects were experienced computer users, but they had little or no prior knowledge in desktop VR, Perseus, or the London collection of Greek vases. The distribution of the subjects was maintained in a round-robin fashion. The test subjects were recruited from the Departments of Computer Science and Electrical and Computer Engineering at Tufts. They were solicited by distributing printed ads and posting notices on course web sites. The subjects were not paid, but they were offered a nominal gift.

5.1.2 Task

The V subjects performed the task of navigating in the 3D Vase Museum (as shown in Figures 4–9) on a laptop PC running the Cortona VRML browser [8]. Figure 9 was the start up screen for this condition, which preloaded all the vases. After a very brief introduction to desktop VR navigation using the cursor keys and the camera views, the V subjects were immersed in an eye-level view similar to Figure 5, from which they would begin working on the task. Using the 3D Vase Museum, the ten task questions could be answered most effectively by either one of the two steps:

1. Use the cursor keys to navigate (walk) in the museum.
2. Select a scene from a camera view (a vase, scatterplot, etc.)

The V subjects were not advised of the benefits of the prototype and were free to explore the features of the 3D Vase Museum on their own.

The P subjects performed the task of browsing the vases in Perseus [9] (as shown in Figure 1–3) on a desktop PC approximately equivalent to the laptop in functional speed and task performance, running Internet Explorer. Figure 1 was the start up screen for the P subjects. They could select a vase page by clicking on its catalog number in the list. The P subjects had full access to Perseus running live on the web. They could search for thumbnail views of the vases (similar to Figure 2), search the vases by keywords or sort the list by other parameters. Any use

of the web that did not pertain to the task was checked. For example, the P subjects were not allowed to abandon Perseus and use another digital library. As long as the P subjects stayed within the scope of the experiment, they were free to use Perseus as they saw fit. Using Perseus, the ten task questions could be answered most effectively by one or more of the following steps:

1. Get the summary information (wares, shapes, etc.) from the collection page shown in Figure 1.
2. Switch to a thumbnail view and get the generic features (i.e. colors, handles, spouts, etc.) shown similar to Figure 2.
3. Sample a specific vase by its catalog number shown similar to Figure 3.
4. Sample a few vases whenever there is a change in the pattern of the catalog numbers (i.e. London A#, B#, etc.)

5.1.3 Procedure

Subjects were given a brief introduction to the system in use and the vase collection. They were given a warm-up question for practice, which did not count toward the final score. Both groups answered the same ten task questions, one at a time, while performing the task. The subjects were asked to perform the task as quickly as possible. Subjects were given as much time as needed. They were not permitted to backtrack once they had completed a question. Subjects were encouraged to pick the best possible answer, rather than leave a question unanswered. The time used to answer each question was recorded by a stopwatch. The subjects were also observed while answering the questions and observations about their use of the system were recorded. At the end of each task question, the subjects filled out the associated task survey, which was not timed. At the end of the session, they filled out the survey and took the quiz, also not timed. The test subjects could skip any survey or quiz questions.

5.2 Results

We found a large and significant difference in speed as well as a smaller but significant difference in accuracy between 3D Vase Museum and the Perseus interface. Subjects using the Vase Museum performed 33% better on the task and equally well on the retention quiz—but achieved this nearly three times faster with our system (13.69 vs. 37.03 minutes).
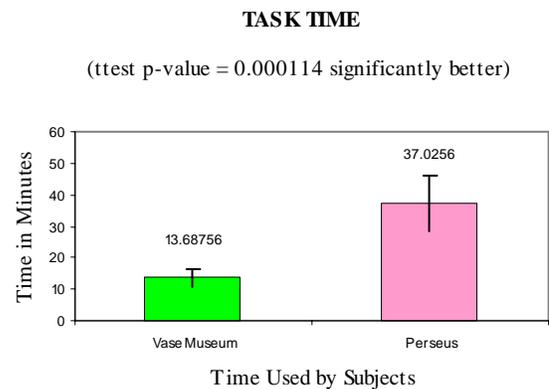
All subjects were able to complete the entire session in approximately 1–1½ hour. Start up time took 5–10 minutes for most subjects. Task time took 30–40 minutes for most subjects. Quiz time took 20–30 minutes for most subjects.

5.2.1 Task Accuracy

The V group performed significantly better than the P group in task accuracy, with $p < 0.0074$, $F(1,18) = 9.11$. The V group scored an average of 7.2 correct answers (out of 10) compared to 5.4 for the P group. The V group scored 5–8 correct answers with a variance of 1.07; whereas the P group scored 4–9 with a variance of 4.8. More details of all the results are available in [19].

5.2.2 Task Time

The V group also performed significantly faster than the P group, with $p < 0.00011$, $F(1,18) = 24.06$. As shown in Figure 12, the V group completed the task in an average of 13.69 minutes, while the P group completed the task in an average of 37.03 minutes.

TASK TIME

(ttest p-value = 0.000114 significantly better)

Figure 12. Time improvement using 3D Vase Museum.

5.2.3 Task Quiz

The V group scored better on the quiz, but not significantly, with $p < 0.375$, $F(1,18) = 0.827$. The V group scored an average of 36.1 correct answers (out of 50 quiz questions), and the P group scored an average of 33.5. We conclude that they showed equal or better retention in this quiz, despite having spent far less time on task.

5.2.4 User Feedback & Observation

In the surveys, the V subjects sampled more vases ($p < 0.036$) in arriving at their answers. They were better able to pick out a unique vase from the collection ($p < 0.038$). They also agreed more strongly with the statements that no reading ($p < 0.0125$) or counting ($p < 0.0351$) was necessary to perform the task.

We also analyzed the subjects' satisfaction questionnaires. They reveal no consistent preference for either system. The V subjects were new to desktop VR systems and paid more attention to the warm-up part of the experiment session. The P subjects appeared to be seasoned web browsers. They glossed over warm-up and exhibited strong individual habits For example, two P subjects preferred not to left-click on the links; instead, they would always right click and select from the dropdown menu. P subjects who were accustomed to searching would go straight to the search field even though the correct information was already in the display. In such cases, their advanced skill set could be an over-kill to the task and served to their detriment.

Each group complained that its system was too slow, but in different ways. The P group was using the actual Perseus system as it appears to users on the web. They were impatient in loading the HTML pages. In trying to answer a collective question, this could be frustrating since it was impossible to view many vases in a single display. Therefore, the P subjects were compelled to open a handful of windows, view the pictures, close them, and open some more. The P subjects 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 subjects were impatient because desktop VR rendering was slow and "walking" could be an ineffective method of access. The system was designed such that the subjects could navigate in the room by "walking" to the vases that interested them. Although this held true during testing, the subjects seemed to

demonstrate a different perception of presence or distance. It appeared that if they were able to see a vase in the scene (i.e. on the screen in desktop VR), they believed they were already there. Therefore, they were impatient in having to hit the cursor keys so many times just to get there or get a closer look. Many subjects would first click on a vase or drag a box around it in an attempt to "walk" to it. For example, a few subjects would click on the vases in the scatterplot (shown in Figure 9) in the hope that they would be immersed in the scene instantaneously.

5.3 Discussion

Our 3D Vase Museum is an example of a solution to the problem of focus-plus-context in information display. In the London vase collection, the focus was given by the vase HTML page and the context was given by the collection. The subjects needed to see both the focus (a vase) and the context (collection) in order to learn what they needed to know for the task questions. The information represented in the vase HTML page (focus) was different from that of the collection (context). The P subjects could choose focus or context, but not both a single display because the topmost window occluded the underlying windows. Given the 1D list (Figure 1) and the individual vase pages (Figure 3), it was difficult to construct a mental picture of the appearance of the collection as a whole. The weak context could further break down when the subjects searched randomly and branched out in links (more contexts).

In our experiment, the P subjects were searching for more pictures of the vases and it appeared that thumbnail views could be more effective for the task. Unfortunately, the thumbnail view (as shown in Figure 2) of the London collection showed all 221 images associated with the 156 vases. Vases with multiple images appeared to be disproportionately representative of the collection while vases with no images would disappear from the display. Nonetheless, thumbnail views could easily be upgraded to offer more image context, if not information context.

The desire for some more meaningful context was evident in our experiment when several P subjects persisted in entering the task questions verbatim in the search field. When the system returned zero search result and offered no contextual cue, some P subjects would rephrase the questions with better keywords.

At the other extreme, some P subjects appeared to have little or no expectation of any context in the current approach. Although the 156 vases could be accessed by scrolling the window shown in Figure 1, many P subjects searched the page habitually, even for a vase already on the screen. For example, they would search for London B 436 (task question #1) without noticing that the vase was already in the collection and on the page. They simply glossed over anything on the screen and searched as many times as they deemed necessary in arriving at their own contextual answers. One P subject attempted to abandon Perseus at start up and requested to perform the task in google.com (request denied).

In our design process, we recognized a deficiency in the way information was stored in the digital libraries. Our 3D world has been meticulously preserved in 1D literature, but this removes information. In other words, describing the 3D Greek vases in 1D text and 2D images (such as what was shown in Figure 3) reduces the information content. For example, a partial description of London B 376 in Figure 6 was:

'*Side A: Two large eyes are depicted, with a nose (or ivy bud?) in between. The center and outlines of the eyes are painted black, the ground is white, and there are purple rings around the pupils.*

Side B: decorated the same as Side A' [9].

Although the words have intrinsic value in literature, it is almost intuitive that subjects who can see a picture of the vase can have a better visual understanding of its appearance. We can claim that the subjects who can visualize and manipulate a 3D model of London B 376 have an advantage over subjects who can only see a 2D picture of it. We can further claim that the subjects who can visualize and interact with a 3D model of London B 376 in a virtual museum have an even greater advantage in learning about it in the collection.

Furthermore, the vases are sorted by catalog numbers, which also loses context. The catalog numbers are merely a vehicle of data management and fail to reflect the archaeological context of the collection. For example, the catalog numbers are sorted in ASCII order, such that vase London B 380 appears before London B 39, encouraging the P subjects to search and in doing so developed an ineffective habit of use over time. This was manifested in the experiment. Although the P subjects spent more time browsing, they did not perform better in the task questions or the quiz.

In light of the above, we have sought to remedy the deficiencies by creating a virtual geometric landscape where the subjects can learn in a natural metaphor using lightweight interactions. We bring focus to highly realistic 3D vases arranged in a meaningful contextual layout. Virtual reality can be a powerful platform for visualizing and interacting with information context because the environment simulates the way we learn about the world around us in real life. The transfer of knowledge can allow us to absorb spatial information with little or no training in a virtual prototype.

6. CONCLUSIONS

In this paper, we demonstrated the feasibility, functionality and usability of the virtual 3D Vase Museum as a new approach to communicating context in a digital library. We highlighted the focus versus context problem and provided a solution using a set of new interaction techniques in 3D focus-plus-context interactive visualization. In the 3D Vase Museum, the user can focus on an individual vase without losing overall context of the collection We conducted an experiment comparing the 3D Vase Museum to its counterpart in Perseus, an established, state-of-the-art digital library, and we found substantial benefits of our new approach.

6.1 Continuing and Future Work

We are currently working to incorporate 3D data models we have built directly into the Perseus Digital Library (the result will be similar to Figure 7). We are also exploring ways of incorporating smaller and less computer-intensive contexts in Perseus. The secondary information space can contain more elaborate world-in-world scenes than text and HTML pages. We have extended the same design principles to other data platforms. For example, we have built analogous 3D prototypes [19] for a Greek coin catalog using two coin collections in Perseus and a London walkthrough created from a set of historical maps [4].

A future implementation of the 3D Vase Museum would be a fully immersive virtual environment, in which a user can browse using innate skills such as turning the head and walking. Our 3D models

are built so that they can be used in such an environment as well as in a VRML browser. The virtual museum can also provide advantages beyond a real museum, like being able to touch or rearrange the exhibits. In an immersive virtual museum, the user can grab a 3D vase with a tracked glove. The user can navigate or fly more naturally by using a 3D input device. The room can be rearranged on command, either by grabbing the vases individually in 3D or by requesting that they be re-sorted along different dimensions. Arrangement and rearrangement in real time is prohibitively slow at present but can be simulated by switching scenes in the current system, though the sense of immersion is reduced during loading. A full-fledged, real-time system for much of our conceptual design will soon be achievable as PC technology improves. However, we hope that the work described in this paper can offer a glimpse of 3D digital library of the future and of its potential benefits [5].

7. ACKNOWLEDGMENTS

8. REFERENCES

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

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

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

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

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

[6] Card, S. K. MacKinlay, J. D. and Shneiderman, B., Editors Readings in Information Visualization: Using Vision to Think, Academic Press, 1999, 307.

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

[8] Cortona VRML Client. Parallel Graphics, http://www.parallelgraphics.com/products/cortona/, 2003.

[9] Crane, G. R. (Ed.). The Perseus Project. http://www.perseus.tufts.edu, Tufts University, Medford, MA., 2003.

[10] Discrete© 3ds max. User References, Version 4, Discrete © Autodesk, Inc. www.discrete.com, 2001.

[11] Giller, V. Tscheligi, M. Schrammel, J. Fröhlich, P. Rabl, B. Kosara, R. Miksch, S. and Hauser, H. 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, 2001.

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

[13] Harward, V. J. From Museum to Monitor: The Visual Exploration of the Ancient World, Academic Computing, 2. (1998), 16-19, 69-71

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

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

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

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

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

[19] Shiaw, H.Y. New Interaction Techniques for the Digital Library: 3D Focus+Context(F+C) Interactive Visualization. Ph.D. Thesis, Tufts University, Medford, MA. 2003.

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

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

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