

Does Zooming Improve Image Browsing?

Tammara T. A. Combs
Computer Science Department
Human-Computer Interaction Lab
University of Maryland
College Park, MD 20742
tcombs@cs.umd.edu

Benjamin B. Bederson
Computer Science Department
Human-Computer Interaction Lab
University of Maryland
College Park, MD 20742
+1 301 405 2764
bederson@cs.umd.edu

ABSTRACT

We describe an image retrieval system we built based on a Zoomable User Interface (ZUI). We also discuss the design, results and analysis of a controlled experiment we performed on the browsing aspects of the system. The experiment resulted in a statistically significant difference in the interaction between number of images (25, 75, 225) and style of browser (2D, ZUI, 3D). The 2D and ZUI browser systems performed equally, and both performed better than the 3D systems.

The image browsers tested during the experiment include Cerious Software's Thumbs Plus, TriVista Technology's Simple LandScape and Photo GoRound, and our Zoomable Image Browser based on Pad++.

Keywords

Evaluation, controlled experiment, image browsers, retrieval systems, real-time computer graphics, Zoomable User Interfaces (ZUIs), multiscale interfaces, Pad++.

INTRODUCTION

In the past two decades, with the emergence of faster computers, the declining cost of memory, the popularity of digital cameras, online archives and even presentation slides, the amount of stored graphical information has skyrocketed. Having the ability to store and manipulate images is becoming more important as images are being incorporated into electronic documents [12]. These digital images are stored and electronically encoded for future retrieval. Hence, there is a growing need for more sophisticated ways of retrieving and browsing images. However, the advances of these tools have not grown as rapidly as the needs of potential users.

There is a vast diversity of users and individual biases that should be taken into consideration as we move toward multimedia systems. Graphical information is being used throughout many systems to help bridge the gap between such differences as languages, gender, age and personality. Sometimes pictures really are worth a thousand words, but what good are they if the interfaces do not offer the support that users need? In this paper we focus on the browsing aspect of the interface.

Browsing is not a new concept. Webster's New World Dictionary gives a basic definition of the term browse, to examine in a casual way. Adults browse for clothes on racks at their favorite department stores and children browse for sweets at their local candy shops. Vendors and department store owners have realized how to capitalize on sales. They know in order to maximize the purchase of their items, browsing needs to be made easy. Most storeowners understand that people will not select what they cannot see. For this reason, merchandise is usually displayed in a manner that best suits the targeted user.

Why should image browsers be any different? Just as librarians shelve books to make them easier for patrons to find, image browsers should display images in such a way that does not distract the user from the main task he/she is trying to perform. For instance, if a user is browsing for an image to include in a document, their browsing experience should not be such that it has made him/her forget the reason they sought the image in the beginning.

In image browsing, screen real estate is very important because it seems as if there is never enough. We believe 3D and zooming make better use of screen space than scrolling. We describe our experiment and give some practical guidelines for future image browsers.

In order to get a basis for understanding the context from which our system was designed, we offer the following definitions:

1. An *Image Retrieval (IR) System* is an application that returns one or more images given some descriptive information. This information can be in the form of:
 - a) An image,
 - b) Keywords or phrases, or
 - c) Natural language.
2. An *Image Browser* is an application that allows users to select one or more images from multiple images. This browser has to:
 - a) Be able to display multiple images at one time (possibly reduced resolution versions), and
 - b) Support inspection of original full resolution versions of an image.

The returned set of images (results) obtained from the Image Retrieval query may be displayed in an Image Browser for further refinement of the search by the user. Often it is the case that the results of the query yield more images than the user desires, so he/she has to browse. It is unfortunate that many query systems ignore browsing and just give the results of the query perhaps in the form of a list. This makes it hard and sometimes impossible for a user to select exactly what he/she needs. After testing the features in many of the browsers we decided to contribute to the image browser community and make our own browser – a zoomable image browser (ZIB) (Figure 1). We designed ZIB where searching and browsing are tightly coupled. With ZIB, the images located in the browse area represent the results of the query posed in the search section. Both the search term and query results can be seen in one view.

To begin our study we evaluated sixteen (16) image browsers (Table 1). We compared and contrasted many features of the commercial and shareware products to discover some of the most popular techniques used in image browsers. We especially targeted software packages that were designed for the purpose of browsing a collection of images. To our surprise, most of the image browsers did not deviate from the typical two-dimensional grid of thumbnails approach. We chose ThumbsPlus (Figure 2) to be the commercial browser we would later use in the experiment because it is a good example of a commercial image browser. ThumbsPlus is a grid of thumbnails that is easy to use and supports access to the full-size image.

Name of Package	Name of Company/Developer
Corel Mosaic	Corel Corporation
PhotoMagic	MicroGrafX Corporation
PhotoCD Access	Eastman Kodak Company
GifDesk	Jay Wherley
Fotoflood Image Manager	EPICAD Design Incorporated
Picture Publisher	MicrografX Incorporated
PhotoDisc	PhotoDisc
Image AXS Pro	Digital Arts and Sciences
ThumbsPlus	Phillip Crews
ACDSee	ACD Systems Unlimited
IrfanView32	Irfan Skiljan
VPIC	Bob Montgomery
CompuPic	Photodex Corporation
Extensis Portfolio	Extensis Corporation
Cumulus	Canto Software
PowerPoint	Microsoft Corporation

Table 1: List of systems analyzed

System Design of the Zoomable Image Browser (ZIB)

We designed a system that integrates image browsing and image retrieval. Query formulation is allowed within the search area. Users have the option of performing a simple or an advanced search. Within the simple search, Users have the option of entering one word or one phrase on which the query will be performed. Within the advanced search, the user may form a query by using a combination of words and/or phrases and boolean connectives. The interface for the search area was written in TCL/TK and the search procedure was written in C++. Once query formulation is complete and the images which satisfy the query have been retrieved, the images are returned within the browse section.

The results of the query appear in the lower (Browse) section and can be browsed by panning and zooming in and out of individual images as well as all images at once. The browse section was built using Pad++, a general purpose engine for writing zoomable user interface [2]. ZIB offers a unique advantage over many browsing systems in that the user has control of the tradeoff between the number of images displayed and the resolution of those images. For example, if ten images are present in the browse section and the user wants to hone in on four of the ten images, he/she can zoom in on the view and see the images enlarge before their eyes. This gives them higher resolution but fewer images. The inverse is also true. Users can zoom out to get lower resolution, but greater numbers of images. Users can also perform in-place zooming which allows them to see an image

at full resolution located in the same place in the same scene.

While users perform successive searches, a history interface maintains a record of previous queries and displays a snapshot of the images that were returned with a particular query. In case users forget the search terms used to retrieve the corresponding set of images, they need only move the mouse cursor over the group of images they wish to inquire about and the search terms will appear at the bottom of the history section. Once users are sure they want a group of images redisplayed in the browse area, they simply click once on the point that is in the history section that depicts the set of images they would like to browse. The user may refer to the history in order to return to a previous search for reevaluation or refinement. The history section was also written in Pad++ and it too allows panning and zooming.

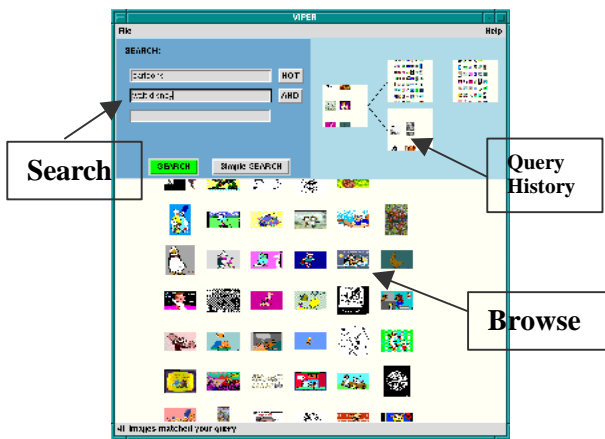


Figure 1: Zoomable Image Browser (ZIB) allows panning and zooming of individual images as well as the entire view.

Search was implemented using standard boolean combinations. This system, however, is not a boolean “free text” retrieval system. Terms permitted within a query are extracted from a thesaurus. The thesaurus defines the controlled vocabulary of allowable terms. The thesaurus was created from the terms used to index the images. These indexed terms were then organized into a taxonomy and presented to the user as the thesaurus. The terms used to index the images are terms that would describe the contents of the images.

Zoomable User Interfaces

Zoomable User Interfaces (ZUIs) are a visualization technique that provides access to spatially organized information. A ZUI lets users zoom in and out, or pan around, to view much more information than can normally fit on a single screen.

ZIB takes advantage of the unique capabilities offered by ZUIs. In a preliminary study, users were able to change their viewpoint and magnification of images within seconds. Many subjects even commented that panning and zooming were very easy and they thought the image-browsing world could benefit from a system such as ZIB.

Although there are many aspects of ZIB that we need to evaluate, we chose to perform a controlled experiment on browsing alone. We were interested in seeing if users would browse differently using 2D, ZUI or 3D environments. The 3D image browsers that we evaluated belong to a company out of Chicago, TriVista, which we are collaborating with. These 3D systems were written in VRML and add a third dimension to image browsing. It was the developers’ hope that the user would view the images as being in a real world environment and thus have an easier and more enjoyable time browsing.

Descriptions of Browsers Chosen

The first image browser we compared to ZIB was ThumbsPlus (T) (see Figure 2), a program that allows users to view thumbnails of images. On the left side of the screen, is a hierarchy of the folders on the current hard drive. The right side is the browsing window, which allows the user to pan via a horizontal scroll bar. Thumbnails can be clicked on, and a full-sized depiction of the image is brought up in a second window. It does not enable users to zoom in or out of the entire view. Because, selected images require a separate window to be opened, there could potentially be $n + 1$ windows, where n is the number of images contained in the current directory.

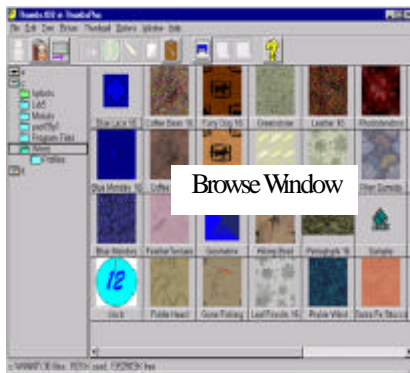


Figure 2: ThumbsPlus (T), a 2D image browser that we used in our experiment.

The second system that we evaluated was the Simple LandScape (L) (Figure 3). Simple LandScape is a 3D image scene that allows you to view billboarded thumbnails of images in 3D as you “walk” or “fly” through the scene. The images will always face you as you move through the scene (using the mouse or the arrow keys). Clicking in a specific thumbnail's area will cause a larger version of the image to appear in a separate window. Users can request a higher overview, or the “from above” view, which enables them to examine more of the images from where they are in the scene.




Figure 3: Simple LandScape (L), a VRML image browser created by TriVista Corporation.

The fourth and final browser we studied was the PhotoGoRound (P) (Figure 4), which draws on a familiar metaphor to organize information: a “lazy-suzan” layout, sometimes used in stores for a card rack or for a jewelry case. The entire scene can be rotated using the upper or lower disk, and clicking the sphere causes the scene to automatically spin. The thumbnails are “billboarded”, so they always face the viewer as the carousel spins. As with the previous two browsers, once the user selects an image, it appears in a separate window.



Figure 4: PhotoGoRound (P), a VRML image browser created by TriVista that uses a “lazy-suzan” metaphor.

Previous Studies

After choosing the systems we wanted to evaluate, we reviewed a number of studies that were similar to the one we wanted to perform. There have not been many studies on how image browser interfaces affect users abilities to select, deselect, or zoom in and out of images. This is especially true when trying to find how different browsers scale up when the number of images is increased.

The first study is the Zoom Browser [10] in which a web-browser (text-only) downloads HyperText Markup Language (HTML) documents from the World Wide Web (WWW) and splits them into thumbnails of pages. Users can navigate through the pages, clicking on links in the text to load new documents. There is also a sense of history keeping in that previously viewed documents remain visible. Users liked the overview achieved from the display of the pages in the Zoom Browser. However, the Zoom browser does not scale up well. Once a certain number of images is displayed on the page, the information displayed on the pages is no longer useful.

In a second study, Protofoil [17], researchers built several information access applications where information (documents and text) was displayed as thumbnails in a grid. Users complained that they were not able to see the contents of the thumbnails clearly so the authors introduced intermediate page sizes allowing users to have a better detail view of the image. There was no concept of zooming used for this image browser.

In a third study, the Pad++ group tested general navigation and history-related effectiveness using PadPrints [9], a WWW companion. PadPrints works

along side a web browser to serve as a history aid building a hierarchy of pages visited by the user. The pages are displayed as thumbnails of images that also serve as links to the represented page. Users are permitted to view the entire graphical history or to zoom in to focus on the particular part of their history. Subjects were asked to navigate the Web with and without the zoomable web companion and for both of the tasks (textual and image-based pages), there were fewer pages accessed, and retrieval time was significantly reduced. This showed that some of the concepts used in PadPrints were effective in navigating. However, PadPrints serves as a web companion and it was not designed as a stand alone image browser. We used some of the same ideas from PadPrints in designing and developing the zoomable image browser. We presume that multiscale contextual display of the images can provide substantial support for browsing.

A fourth, and particularly relevant study [11], is that of a group of students from the University of Maryland. The main focus of this study was to come up with an optimal tradeoff between image size and the number of images that could be displayed at once. They found that increasing the number of images while reducing their size resulted in reduced task completion time. However, they only tested a maximum of thirty-six (36) images. They concluded that further testing should be done with larger image sets, to determine what the optimal number of images viewed simultaneously are.

EXPERIMENT

We performed a user study to assess each of the browsing systems. We adopted the hypothesis that there would be no statistically significant differences in the time it took users to locate the targeted images, the browser users preferred, or in the number of incorrect selections made on a particular browser. This user study however did result in a statistical significance in each of these dependent variables, as we will discuss below.

Equipment

ThumbsPlus, Simple LandScape and PhotoGoRound are windows-based programs while the zoomable image browser runs on Linux. Because each subject evaluated all four browsers, we used two machines to avoid switching between operating systems on one machine. We were careful to eliminate any windows

management tasks to avoid any differences in the two operating systems.

Both of the computer systems used were 166 MHz Pentium PCs with 17" monitors. One system was running Windows NT 4.0 with a resolution of 1024 x 768, while the other system ran Linux with a resolution of 1280 x 1024. Because we wanted the machines to be of comparable speed, the Windows NT machine had 114 megabytes of RAM and the Linux machine had 64 megabytes of RAM. Browsing with the windows-based systems was performed using a 2-button mouse and browsing with the zoomable browser required the use of a 3-button mouse. We wrote a program that automated and recorded the questions and tasks presented to the subjects.

Stimuli

Each subject was asked to browse through a set of images until he/she had located the target image. They used browsing functionality specific to the browsing system they were working on at the time. For example, a user may have enabled autospin in the PhotoGoRound to complete the task of finding an image of strawberries as pictured below in Figure 5.



Figure 5: Example image that subjects were asked to browse for during the experiment.

Training

Before beginning the experiment, each participant was educated in the use of image browsers in general. We wanted to be sure each subject had a clear understanding of the assignment they were about to perform. Subjects completed five pre-tasks using the first image browser that they would use. There was no time limit to the pre-tasks and subjects were

informed that they did not have to proceed with any further tasks until they felt comfortable with the browsing system. The goals of training were to verify that our subjects understood the navigation techniques for the browser and also to familiarize them with the program we used to automate the questions.

Because training had already been administered to subjects on browsing and performance times were not being recorded, we did not repeat all of the previous techniques for training on the secondary browsers. Subjects were simply given the sheet of instructions before beginning the tasks associated with that browser and asked to read it and continue with the experiment when they were ready.

Method

The primary design of this experiment was a 3x4 block design. Each subject was given matched tasks for the four browsers. The independent variables were the different browsers (ThumbsPlus, Zoomable Image Browser, Simple LandScape, and PhotoGoRound) and number of images (25, 75, 225). The orderings of both independent variables were randomized.

We conducted two experiments simultaneously. The first experiment used the method of between-subject testing. Each participant was randomly assigned one of the four browsers as their primary browser for use in this first experiment. Each user was instructed to browse through each of the three image sets. When the test user was ready, he/she would initiate a request for a new image. They would then return to the browser assigned to then and signify that they had found the correct image by selecting it. Browsing time for each of these images was recorded.

The user repeated this process until the five images for each image set had been correctly selected. After each image set of five images, there was a small task to measure recall. Subjects were instructed to find an image they would like to send to a friend. Time was not recorded for this task. The purpose was to give subjects exploratory time as well as to observe how they would use the browser when there were no time constraints. Subjects were presented with four images to test recall. Subjects were instructed to select either "Yes, the image was in the set of images" or "No, I don't recall seeing this particular image in the set." They then evaluated their primary browser using a questionnaire and they were asked to give feedback on anything they felt wasn't addressed during the

experiment and they were also asked for any suggestions for improvement of the system.

After the debriefing segment of each test user's primary browser, they each began the second experiment, a with-in-subject test. With-in-subject design requires all the participants to use all of the systems that are being tested. Participants evaluated their secondary browsers in random order. In addition to the browsers being presented in random order, each participant only evaluated one of the three image sets. For example, one test user evaluated ZIB as her primary browser using the set of 25 images, then with 75 images, and then with 225 images. She then evaluated Simple LandScape with 25 images, PhotoGoRound with 25 images and ThumbsPlus with 25 images. Just as the order for evaluation of the three image browsing systems was randomized, so were the image sets.

Sixteen dependent variables were analyzed in the experiment. Mean performance time was measured for each of the three sets of images for the primary browser. Time to complete the task was calculated from the time that the subject initiated a request for a new question until they completed the task by selecting the targeted image. The number of incorrect selections was measured for the three image sets for the primary browser and the three secondary browsers. If an incorrect selection was made, subjects were instructed to continue searching until the correct image was found. A correct selection was eventually made 100% of the time. Percentage correctly recalled was measured for each of the three image sets for the primary browser. We calculated this by placing the total number correctly recalled over four (total number possibly correct).

Lastly we measured mean subjective satisfaction ratings for each browser. We calculated these using questions from the Questionnaire for User Interaction Satisfaction (QUIS) developed at the University of Maryland [20] as well as questions specifically related to image browsing. All of the questions were based on the QUIS format and were therefore on a scale of one to nine.

Subjects

There were 30 participants involved in this experiment, most of whom were students at the University of Maryland, College Park with various backgrounds including Computer Science (45%), Electrical Engineering (20%), Graphic Design (10%)

and Library Information Services (20%). Approximately 40% of the subjects were female and 60% of the subjects were male.

Participants' ages were recorded using ranges so they would not feel uncomfortable disclosing their ages. From the data we collected, 45% of the participants were between the ages of 18 and 25, 40% between 26 and 35 and 15% between 36 and 45. 97% of subjects reported they were experts on the World Wide Web (WWW) with the average user browsing 14 hours per week. Users also reported using a personal computer (PC) an average of 36 hours per week.

Each subject was paid \$10 for participating in the experiment.

Results

We observed a statistically significant interaction effect between the browser and the number of images viewed with that browser for performance time. ZIB proved to be faster than the other browsers for each image set, although it was only significantly faster than Simple LandScape and PhotoGoRound ($F_{2,18} = 12.359, p < .0005$). Even with 225 images, ZIB was not significantly faster than ThumbsPlus (See Figure 6). For an effect to be considered significant, p had to be less than or equal to 0.05.




Figure 6: Mean performance time per browser for each image set.

There was also a statistically significant interaction effect between browser and number of images for user satisfaction (see Figure 7). Again, ZIB had the highest user satisfaction ratings, but it was only significantly faster than Simple LandScape and PhotoGoRound ($F_{3,29} = 15.667, p < .0005$). ZIB was not significantly preferred over ThumbsPlus.




Figure 7: Mean user satisfaction ratings for all browsers.

There were no significant ordering effect as it relates to user satisfaction for the primary browser ($F_{3,25} = .745, p = .535$). Also there were no significant differences in the browser test subjects preferred for the set of 75 images ($F_{3,20} = 2.463, p = .092$) or 225

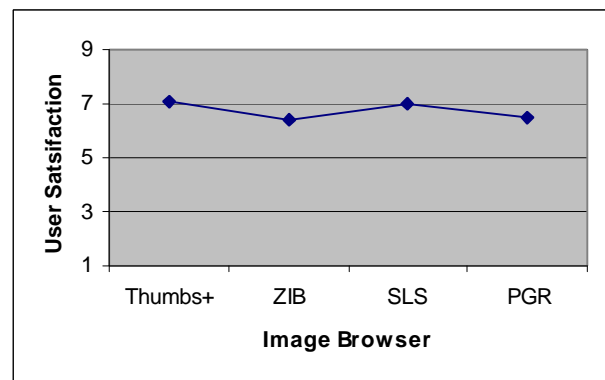


Figure 8: Mean user satisfaction ratings for primary browsers averaged on all image sets.

images ($F_{3,26} = 2.127, p = .121$).

Subjects recalled images with 15% higher accuracy using ZIB compared with ThumbsPlus with our maximum number of images (see Figure 9).




Figure 9: Average percentage of images correctly recalled.

Also, there was a large number of incorrect selections (Figure 10) made for PhotoGoRound and Simple LandScape while relatively few were made for ThumbsPlus and ZIB.




Figure 10: Total number of incorrect selections made.

Analysis

The results of this experiment showed that the zoomable image browser as well as the traditional 2D grid of thumbnails works best for performance time and user satisfaction. Users also made fewer incorrect selections for ZIB and ThumbsPlus. While the above statements are certainly true, we should note that all browsers did fairly well with performance time and recall with the small image set. With the maximum number of images, there was no preference toward ZIB, but ZIB had the fastest performance time.

Another peculiar observation we made was that roughly half of the subjects did not zoom when given 225 images in ZIB despite the fact that we gave training in zooming. However, there was still a performance time improvement. Perhaps it was

because all the images were on one screen and they never had to adjust the view if they chose not to.

We decided to maximize all browser screens to give the user maximum browsing space. However, this introduced a confounding variable because the 3D browsing systems used less screen space than the other two browsers did. This was due to the setup and design of the 3D systems, which was out of our control. Perhaps this is why performance time for the 3D systems was not as fast as the other browsers.

Recall

Test users had to do a substantial amount of scrolling in ThumbsPlus with 225 images. Perhaps this accounts for the 15% difference in recall compared to ZIB. Conceivably moving the scroll bar distracted them from the task and they were unable to remember the images that they had just stored in their short-term memory.

Incorrect Selections

A selection was considered incorrect when the user selected an image other than the target. Once an incorrect selection was made, the user continued to browse until the correct image was selected. Most incorrect selections were accumulated with the PhotoGoRound and Simple LandScape browsers. Perhaps this is due to the movement of the scenes. Users tried to select images from the PhotoGoRound while it was still spinning. Most of the time the result was the selection of an unwanted image. On the other hand, incorrect selections were relatively low in ThumbsPlus and ZIB. An observation that we made, as it relates to ZIB, is that as the number of images increased, so did the number of incorrect selections. Oddly enough, this was the only browser where there was a direct correlation between number of images and number of incorrect selections. An explanation we offer for this was gathered from observing the subjects. Despite having 225 images on the screen at one time, most users still did not zoom. They stayed zoomed out and thus could not see the images clearly.

Qualitative Results

We gathered some qualitative results from our users as they performed the experiment. While many subjects said PhotoGoRound was the most entertaining, the most popular comment was that users did not like or wanted to change the speed of rotation. The Zoomable Image Browser, repeatedly said to be the easiest, received many comments

suggesting the ability to group images in clusters by content. ThumbsPlus also received requests for an added vertical scroll bar, more accessible zooming, more images per page and the disappearance of the explorer window once their image set had been selected. The most sought after feature subjects wanted added to the Simple LandScape had to do with the overview. Users wanted some way to globally view places they had already visited in the landscape. Moreover, they wanted to see where they were presently in relation to the entire plane. Subjects repeatedly stated they were lost.

We purposely left out searching tasks in this experiment. However, many subjects explicitly expressed a desire to search for the target image rather than browse for it.

CONCLUSION

While the current study shows some preliminary results, there are still several unanswered questions. For one, is there an optimal number of images that should be displayed on a screen at one time? If so, at what resolution should they be viewed? At what point will users feel a need to zoom in or out of their current view? Perhaps we should have had a fourth image set of 500 or more images, so that users would have had to zoom in order to see the contents of the images. Or perhaps we should have used a smaller window for the same reason.

There are many unanswered questions, but from this experiment we have come up with some practical guidelines for designers of image browsing systems. Designers should choose approaches such as a zoomable image browser or 2D grid of thumbnails if they are concerned about the number of incorrect selections users make. The number of images displayed in the browser is also important. We saw in the results that there was a significant interaction effect between browsers and number of images. This means that designers should decide if their image browser is going to be used for large or small image sets. Either of the four aforementioned browsers are fine for relatively small numbers of images, but more traditional approaches or our zoomable image browser appear to work better when there is a large number of images.

ACKNOWLEDGEMENTS

We would first like to thank John Mareda and Ed Marek of TriVista for their input and support throughout this work. We give a special thanks to

Norina Dixon who helped in the original implementation of ZIB. We thank all of our test users who so graciously participated in our experiment. We say thank you to the members of HCIL at UMD for their useful comments and accommodating spirits and also John Jones and Maya Venkatraman who helped with the analysis of the data. In addition, we acknowledge UNM and other members of the Pad++ team, especially Jim Hollan and Jon Meyer.

The Zoomable Image Browser and Pad++ in general have been largely funded by DARPA to whom we are grateful. This study was funded by TriVista Corporation.

REFERENCES

1. Bederson, Benjamin, Hollan, James, "Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics", *Proceedings of UIST'94 ACM Press*, 17 – 26.
2. Bederson, Benjamin, Hollan, James, Perlin, Ken, Meyer, Jon, Bacon David, and Furnas, George, "Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics", *Journal of Visual Languages and Computing* (7), 1996, 3 – 31.
3. Besser, Howard, "Getting the Picture on Image Databases", *Database v. 18, April/May 1995*, pp. 12-19.
4. Chapman, Stephen, "Questions on Image Browsers", *Database v. 18, Oct/Nov 1995*, p. 6+.
5. Eakins, J., "Pictorial Information Systems - Prospects and Problems", *14th Information Retrieval Colloquium, April 1992*, pp. 102-123.
6. Edelstein, Herbert, "Document Image Management", *DBMS, April 1992*, pp. 46-52.
7. Gerfelder, Norbert, Kromaker, Detlef, "Harmonization In the Area of Image Communication", <http://icib.igd.fhg.de/icib/descript/topics.html>
8. Gudivada, Venkat, Raghavan, Vijay, "Content-Based Image Retrieval Systems", *Computer, IEEE, September 1995*, pp.18-21.
9. Hightower, R., Ring L., Helfman, J., Bederson, B., & Hollan, J. (1998), "Graphical Multiscale Web Histories: A Study of PadPrints", *In*

Proceedings of ACM Conference on Hypertext (Hypertext 98) ACM Press, pp. 58-65.

10. Holmquist, Lars Erik, "Focus + Context Visualization with Flip Zooming and the Zoom Browser", *Proceedings of ACM CHI '97 Conference on Human Factors in Computing Systems Extended Abstracts*, 1997 v.2, p. 263-264.
11. Liebeskind, Andrew, North, Christopher, Orandi, Sharam, "Thumbnailing Techniques: Browsing Large Image Databases", University of Maryland, College Park, December 1994
<http://www.cs.umd.edu/users/north/thumbnail.ps.Z>.
12. Korfhage, Robert, *Information Storage and Retrieval*, Wiley Computer Publishing, John Wiley & Sons, Inc., New York, 1997.
13. "The Kodak Photo CD: Photo CD Application background", <http://www.itc.nl/~bakker/info/rs-data/kodak.htm>
14. Ousterhout, John, *Tcl and the Tk Toolkit*, Addison-Wesley, 1994.
15. Plaisant, Catherine, Carr, David, Shneiderman, Ben, "Image-Browser Taxonomy and Guidelines for Designers", *IEEE Software*, vol. 12, no. 2, pp. 21-32.
16. Rao, Ramana, Card, Stuart, Jelinek, Herbert, Mackinlay, Jock, Robertson, George, "The Information Grid: A Framework for Information Retrieval and Retrieval-Centered Applications", *UIST '92, November 1992*, pp. 23-32.
17. Rao, Ramana, Card, Stuart, Johnson, Walter, Klotz, Leigh, Trigg, Randall, "Protfolio: Storing and Finding the Information Worker's Paper Documents in an Electronic File Cabinet", *CHI '94*, pp.180-185.
18. Sclaroff, Stan, Taycher, Leonid, La Cascia, Marco, "Image Rover: A Content-based Image Browser for the World Wide Web", Boston University *Technical Report #97-005*, March 1997.
19. Shneiderman, B. (1997). *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Addison-Wesley.
20. Slaughter, L., Harper, B., & Norman, K. (1994). Assessing the Equivalence of Paper and on-Line Versions of the QUI5 5.5. *In Proceedings of 2nd Annual Mid-Atlantic Human Factors Conference* pp. 87-91.
21. Spense, Robert, "A Taxonomy of Graphical Presentation", *ACM, INTERCHI '93 Adjunct Proceedings*, pp. 113-114.
22. Ventura, Della, Gagliardi, I., Schettini, R., Solari, R., "An Iconic Browser for Image Databases", *Proceedings of the 7th Intl. Conference of Image Analysis and Processing III*, September 1993, pp. 279-286.