THREE-DIMENSIONAL INFORMATION SPACE: AN EXPLORATION OF A WORLD WIDE WEB-BASED, THREE-DIMENSIONAL, HIERARCHICAL INFORMATION RETRIEVAL INTERFACE USING VIRTUAL REALITY MODELING LANGUAGE

DISSERTATION

Presented to the Graduate Council of the
University of North Texas in Partial
Fulfillment of the Requirements

For the Degree of

DOCTOR OF PHILOSOPHY

By

Peter Scannell, B.S, M.Ed.
Denton, Texas

December 1997

TABLE OF CONTENTS

LIST OF TABLES

CHAPTER 1

INTRODUCTION

The World Wide Web (WWW), first conceived in March of 1989 by Tim Berners-Lee and officially released to the Internet community at large in the summer of 1991 [Berners-Lee, 1996 #31; Krol, 1996 #9; Hughes, 1994 #7], has now gained a significant presence on the Internet in terms of user access to all aspects of the Internet. The Internet is a large collection of computers, each of which is connected to the others and each of which contains information. The World Wide Web provides a graphical user interface to the information available on the Internet. As Pesce [, 1995 #10] has written, the World Wide Web is “the space consisting of all network accessible information” (p. xxi), a description formerly applied to the Internet itself. Unfortunately, network accessible does not necessarily mean human accessible or even human findable. Janes and Rosenfeld [, 1996 #8] complained that, with the increasing amount of information available through the World Wide Web, “it has become increasingly difficult to find information of value, or information relevant to a particular subject” (p. 711), even though they listed a number of ways that have been created to find information on the World Wide Web. The systems that Janes and Rosenfeld described concentrate on organizing information and resources into hierarchical categories or indexed databases, with little emphasis on providing graphical search interfaces to the information.

Although a number of different types of search interfaces to information databases have been created in recent years, the most common interface on the World Wide Web is still the simple entry field for keyword searching or hierarchical, text-based menu systems. Two-dimensional (2-D) and three-dimensional (3-D) search interfaces have not yet had any significant presence on the World Wide Web. Recently, however, interest in three-dimensional, artificial worlds, or virtual reality, "has gone from an obscure scientific toy to being touted as the future of computing" [Pimental, 1993 #2, p. xiii]. Interest in three-dimensional Web applications, including three-dimensional search interfaces, is also increasing, with the development and deployment of the Virtual Reality Modeling Language (VRML).

Although the idea of VRML was first proposed in May 1994 [Pesce, 1995 #10, p. 23] and is still being extended, it is now the defacto standard for World Wide Web three-dimensional development [Broll, 1996 #6]. Currently, VRML provides for the creation of static 3-D scenes with limited interaction possibilities, but VRML viewers already allow the user "to navigate freely through the scene and to follow hyperlinks to new 3-D worlds, HTML documents or other valid MIME types" [Broll, 1996 #6, p. 427]. HTML stands for Hypertext Markup Language, and MIME stands for Multipurpose Internet Mail Extensions. Even though 3-D interfaces on the Web are new, they are already being presented as the answer to our problems of finding information on the Web [Pesce, 1995 #10]. Summitt and Summitt [, 1995 #16] described the advantages of 3-D:

> If we can find a way to "sense" this two-dimensional, virtual location, then the location becomes "real" in our perceptions. This way of "sensing" must take place in a three-dimensional reality because we human beings are born into a three-dimensional navigational space in life. (p. 168)

Overview of the Problem

With so much being written about the World Wide Web generally and about VRML in particular, it is to be expected that developers will rush in to provide 3-D applications. Among these will undoubtedly be 3-D search interfaces and virtual worlds to access data on the Web.

Currently, most VRML applications are little more than experiments to see what can be built. Part of the reason why this is so is that the language specification is still very new and is still undergoing development. Another problem with developing VRML applications is that powerful machines (minimum Pentium 100 with 32 MBytes of memory) and high network bandwidth (minimum 28.8 modem) [Flohr, 1996 #13] are required to run them with any reasonable degree of speed. Despite this, applications are being built to aid in visualizing architectural and industrial models, to provide aid in navigation around Web sites [Flohr, 1996 #13], and to provide pictures of geographical locations [Broll, 1996 #6; Pesce, 1995 #10].

The experience of the visualization community with applications written for visualization tasks gives an indication of how much testing VRML applications will undergo as they are written. As Globus and Uselton [, 1995 #3], writing about the state of visualization software, observed:

> The grand hypothesis of the visualization community is that scientific visualization improves human insight. Several methods have been traditionally used to prove this hypothesis:
>
> • Proof by repeated assertion.
>
> • Proof by vigorous gesticulation.
>
> • Proof by pretty picture. (p. 43)

In other words, most applications in the visualization community are deployed without any testing to determine how well users can use them or whether they are better or worse than current scientific visualization tools. Unfortunately, the same state of affairs seems to be true for 3-D and virtual reality applications currently available on the Web. Many 3-D applications are being built and deployed, with little more than pretty pictures to justify their worth.

Purpose of the Study

This research examined the differences between a 3-D search interface to an information database and a conventional, text-based, one-dimensional (1-D) interface to the same database on the World Wide Web. The intent of the research was to determine how well a 3-D search interface worked in comparison to a 1-D search interface on the World Wide Web.

The research focused on how well users could find information using the 3-D search interface versus the 1-D search interface. In particular, claims have been made that 3-D search interfaces are more natural than 1-D search interfaces because they correspond more closely to the real world and should enable users to work more quickly and with greater satisfaction compared to 1-D search interfaces. On the other hand, 1-D search interfaces are generally easier to use and run faster on standard computers than 3-D search interfaces. In addition, users would probably be more familiar with 1-D search interfaces. This study hypothesized that there would be a difference in users' ability to find information using the 3-D search interface versus the 1-D search interface and in the attitudes towards the interfaces.

A number of World Wide Web information search systems contain hierarchical information; however, for the purpose of this study, a 3-D search interface to a section of the Yahoo search system was created. The Yahoo system provides an interface that allows users to search for information using a hierarchical menu system to navigate ever deeper in the menu structure until either the desired information is found or the search reaches a dead end. VRML was used to build a three-dimensional representation of the menu system, with 3-D objects representing each of the nodes in the menu system and other 3-D objects representing each of the referenced Web pages. The 3-D search interface, created for this study, built is similar to the Cone Tree interface [Robertson, 1991 #63], in which a three-dimensional picture of a file system is created by building nodes to represent each directory and file in the file system. The VRML interface is different in that Cone Trees provide for only a single point of reference into the system, whereas the VRML interface allows users to move into and through the nodes of the system. Cone Trees and their derivatives are well suited to model a hierarchical system of files or menus.

The 3-D search interface was tested against the 1-D search interface to determine how well users interacted with the 3-D search interface versus the 1-D search interface. This helped to determine whether it is worth the time and effort required to build a 3-D search interface to a search engine, rather than the much simpler 1-D search interface that current menu-based systems use.

Hypotheses

1. Users will take differing amounts of time to find information using the 3-D search interface and the 1-D search interface.

2. User satisfaction after using both search interfaces will be different for the 3-D search interface and the 1-D search interface.

3. Users will prefer either the 3-D search interface or the 1-D search interface.

4. Users will rate the 3-D search interface and the 1-D search interface differently.

Significance of the Study

This research is important because there is a general trend in Web development towards interfaces that are "cool," with little thought as to how usable they are. As more and more people assert that 3-D search interfaces are inherently better than 1-D search interfaces, it is only a matter of time before a 3-D search interface will be built and deployed on the Web. Before this happens, 3-D search interfaces must be tested to determine the general usability of the interface, the efficiency and speed of the interface, and the user's satisfaction with such interfaces.

The testing of a 3-D search interface for information retrieval, when compared to the tested 1-D search interface, provided information on whether 3-D search interfaces improve users' attitudes and performance when searching for information on the Web. The comparison of the 3-D search interface to a standard 1-D Web search interface, provided an indication of how well a user did with the 3-D search interface versus the standard search interface and how satisfied the user was with the 3-D search interface versus the standard search interface. These results also provided helpful information for those considering developing a 3-D search interface for information retrieval.

Limitations of the Study

The study is limited by the following:

Version 1.0 of VRML was used in the study. It has limited capability for user interaction and does not have the ability to provide animation. Future versions of VRML will have more capabilities for both user interaction and animation.

The research was conducted using Pentium 166 computers with 32 MBytes of memory and 3-D graphics cards. The results may not be applicable to slower systems because interaction with the 3-D search interface may be too slow for users to tolerate and could significantly reduce user satisfaction with the 3-D system. The results also may not be applicable to faster systems because part of the user satisfaction with the system may be dependent on the speed of interaction. A faster system could increase user satisfaction simply because of the increased speed of interaction.

The search interface used the Yahoo search interface for comparison searches. Other systems use different hierarchical organization schemes. Extrapolating users' interaction with the Yahoo hierarchical system to other systems may not be possible because the structure of the hierarchical system used might have a significant effect on the results of the study.

The sample population consisted of undergraduate and graduate students in Computer Education and Cognitive Systems classes, and graduate students in the School of Library and Information Science at the University of North Texas. The study results may not be valid for other populations.

The program used to run the VRML application was Netscape Navigator 3.0 with the Live 3D plugin. Netscape Navigator 3.0 with the Live 3D plugin has a specific mode of interaction with three-dimensional objects and with the three-dimensional world that it displays. User interaction is different for each 3-D browser, and results obtained with the Netscape Navigator browser may differ from those obtained with another 3-D browser.

The number of data points tested was limited to those under the Education category in Yahoo. The results obtained with this sample of data points may not be applicable in a situation in which there are many more data points, with a resulting increase in the complexity of the 3-D search interface.

Definition of Terms

Alta Vista--One of a number of Internet-based search engines that provide an index to documents on the World Wide Web.

browser--The software that displays HTML documents which have been downloaded from the World Wide Web.

hypertext--A system of linking documents together, such that any document can contain reference links to other documents. The reader of the document can use these reference links to open the referenced document by selecting the reference link.

HyperText Markup Language (HTML)--The tag-based language that determines the formatting of documents on the World Wide Web.

Internet--The global network of computer networks.

Live 3D plugin--The plugin that enables the Netscape Navigator browser to display 3-D images created using VRML.

Lycos--One of a number of Internet-based search engines that provide an index to documents on the World Wide Web.

Multipurpose Internet Mail Extensions (MIME)--An extension of the Internet mail standard which is used to package and identify multimedia resources transmitted over the Internet and World Wide Web.

Netscape Navigator browser--The browser developed by the Netscape corporation.

network--A term for physically separated computers connected to each other by means of telecommunications hardware and software.

one-dimensional search interface (1-D search interface)--A document search interface that only displays search items via a simple menu or list of items.

plugin--A piece of software that extends the capabilities of a browser.

scientific visualization--The process that takes numerical data and constructs meaningful images from it.

three-dimensional search interface (3-D search interface)--A document search interface that displays search items in a three-dimensional coordinate system.

two-dimensional search interface (2-D search interface)--A document search interface that displays search items in a two-dimensional coordinate system.

Virtual Reality Modeling Language (VRML)--The programming language that describes three-dimensional objects and environments in such a way that standard World Wide Web browsers can display them. Files can be small and are delivered over the World Wide Web.

visualization software--Software that takes numerical data and produces meaningful images.

Web--A short name for the World Wide Web.

Web Crawler--One of a number of Internet-based search engines that provide an index to documents on the World Wide Web.

World Wide Web (WWW)--A global hypertext system that runs on the Internet and provides links to documents on the Internet using HTML and a system of addressing that pinpoints a specific document.

Yahoo--One of a number of Internet-based search engines that provide an index to documents on the World Wide Web.

CHAPTER 2

CURRENT RESEARCH

Graphical Information Retrieval Systems

.c3.Introduction

Information retrieval systems were originally developed from traditional retrieval systems such as the card catalog. In addition, most current systems were initially developed when text-based terminals were the standard. In practice, this is still the standard in many facilities, with graphical workstations only now replacing text terminals. As a result, most current information retrieval systems are text-based, using keywords to probe the database and return a list of hits on the database. These information retrieval systems suffer from a number of problems.

Korfhage [, 1991 #50] provided the following list of potential pitfalls with traditional systems:

> They present to the user a limited set of documents, those judged most relevant. The judgment, however, is that of the system, not of the user.
>
> The retrieved documents are presented in a sequential list, often with little or no indication of the system's evaluation of each document. At best, the list is in decreasing order in terms of the single measure of relevance used by the system.
>
> There is no clear indication of the relationships that may exist among the retrieved documents.
>
> The user has little or no idea of the precise criteria for retrieval or nonretrieval of a document.

What is presented to the user is often only a bibliographic record, not a full document.

The user must assume that the system has made appropriate choices, and has little or no information about possible relevant documents that the system may have rejected.

The user is treated as a uniform, faceless entity: there is little or no effort to take into account individual user differences.

In many systems the user is left with little or no guidance in reformulating a query, should that be necessary.

Even in systems with relevance feedback, the input that the user has in the feedback process is limited, and must be given with little guidance from the system. (p. 134)

Even when the user has retrieved a list of documents, it is difficult to tell just what has been obtained. Many questions can be asked about the list of documents returned. Has every relevant document been returned? What else is out there? What is the relationship of these documents to each other? What is the relationship of these documents to other documents not retrieved? If one of the documents returned is judged to be especially relevant, how can other related documents be identified?

Most current systems do not provide adequate answers to these questions. As information databases grow in size and the user's expectations of information retrieval systems increase, our systems must do a better job of information retrieval. In addition, computing power is increasing at such a rate that strategies for information retrieval that would have been prohibitively expensive in terms of computing power and time are now becoming practical.

Many researchers are working on perfecting schemes that determine the user's exact search requirements and that will return all documents and only those documents relevant to the user's search request. Other researchers are working on schemes that will allow the searcher to deal easily with hundreds or thousands of documents at once and to select relevant documents based on cues provided in the search interface. Graphical retrieval systems seek to increase the number of documents that can be displayed simultaneously and to allow the user to select relevant documents based on visual cues in the interface.

<u>The Graphical Solution</u>

In order to solve the problems inherent in current retrieval systems, Korfhage [, 1991 #50] suggested that "the viewpoint [of information retrieval systems] should shift from retrieval to display" (p. 134). Fowler, Fowler, and Wilson [, 1991 #56] wrote that "one of the common goals of visualization is to help manage and understand large amounts of data or information. From this perspective, information retrieval systems are natural candidates for visualization techniques" (p. 142).

Instead of concentrating on returning just the set of documents that the user deems most relevant, systems should concentrate on displaying the documents in such a way that the relationship between documents returned (and, if possible, those not returned) can be easily determined. The user can then use iterative refinement processes to home in on the set of most relevant documents. Instead of the system doing all the work, visual and graphical representations allow the user's own visual processing abilities to play a key role in the retrieval process.

Olsen, Korfhage, Sochats, Spring, and Williams [, 1993 #55] described some desirable features of graphical information retrieval systems:

> All documents should be presented via graphical representations in one display, thus avoiding the problem of forcing the user to create a holistic view from different displays.
>
> The position and other graphical features of a document's icon should intuitively give information on a document.
>
> Data reduction may be necessary, but the most important document attributes, as defined by the user should be retained in the display.
>
> The display should give an overview of the complete document collection retrieved, as seen from the user's perspective.
>
> Users should be able to identify single documents for retrieval of additional information.
>
> Users should be able to change the display interactively by relating document attributes to new graphical features, by viewing documents from a new perspective, etc. The transition from one perspective to the other should be apparent. (p. 73)

Graphical solutions can be one-dimensional, two-dimensional, or three-dimensional. One-dimensional systems would consist of a list of icons or graphics, but would not differ significantly from text-based information retrieval systems because the text list is merely replaced by icon or graphics lists.

Two-dimensional systems provide the potential for a geometrical increase in information, because the added dimension allows for the possibility of multiplying the information from both dimensions together. Three-dimensional systems would then potentially provide even more information, with the information from three dimensions multiplied together.

Ware, Arthur and Booth [, 1993 #67], on the other hand, argued that 3-D systems will probably not result in quite the information gains that we might expect. As Ware et al. observed,

> Using a naive view, moving from 2-D to a 3-D display should vastly increase the amount of information that can be represented. Consider a 1000 x 1000 computer display. On a line we can perceive 1000 distinct pixels; on the plane we can perceive 1000^2 distinct pixels. Extending this logic we should be able to display 1000^3 distinct voxels in a 3-D volume.
>
> Clearly, there is a flaw in this logic; in general we do not perceive volumes of data; we do not perceive details of the insides of solids, only the layout of surfaces in space. This places an upper bound on the amount of information that can be represented in a 3-D space.
>
> Using a different naive view, we can argue that because stereo 3-D viewing is based on input from two eyes this can at most double the amount of information perceivable, and we know that such effects as binocular rivalry mean that truly independent images from the two eyes cannot be perceived.
>
> This second view suggests that moving from 2-D to 3-D will yield only a small benefit in visualization, at least for the understanding of abstract data.
>
> There is a third view based on the ecological argument that because we have evolved in a 3-D world, information presented in 3-D will be processed more easily by the visual system. Networks of information do not have an inherent dimensionality in the geometric sense, but if our brains prefer 3-D layouts, then a 3-D layout may be more effective in conveying the information. (p. 123)

This third view is the one adopted by advocates of virtual reality when extolling the virtues of 3-D systems.

Many experimental systems have already been built to test the advantages of 2-D and 3-D graphical information retrieval systems. The following sections describe some of the systems that have been developed for both 2-D and 3-D viewing. The sections have been divided into the systems that present the data in a two-dimensional display and systems that present data in a three-dimensional display.

Current Two-Dimensional Information Retrieval Systems

.c4.The Information Grid. Rao, Card, Jellinek, Mackinlay and Robertson [, 1992 #66] described a system for searching a document space where images of the documents are stored. The system allows the user to enter search terms in a query area. Retrieved documents are then shown in a grid of fixed-size thumbnail sketches, where thumbnails are generated from the first page of each document. Thumbnail documents can then be used to seed further searches by selecting suitable documents from the results area or from a holding area where documents have been dragged for storage.

.c4.Protofoil. Rao, Card, Johnson, Klotz and Trigg [, 1994 #61] built on the work described in the Information Grid to handle paper documents. It stores, retrieves, and manipulates paper documents as electronic images. The four components of the system include the following:

> a document database for storing documents as page images along with other information needed to support search, retrieval, and browsing of the documents.
> document services for printing, e-mailing, faxing, and OCR.
> scanstation user interface for storing documents into the document database and for otherwise accessing documents or invoking services using paper user interfaces.
> workstation user interface for retrieving and browsing documents and auxiliary information and invoking services. [Rao, 1994 #61, p. 182]

The Protofoil system provides document retrieval "based on an iterative loop in which four user actions are variously interleaved: selecting a scope (initially the entire file cabinet); specifying a query; browsing results; and browsing and using documents." (p. 182) Searches return results in different five views: thumbnails--array of small pictures of the documents returned; description--a list showing attributes of the documents returned; category groups--grouping of documents according to their category; clusters--list of automatically generated clusters of documents, with each cluster showing the number of documents in the cluster, the central terms of the cluster, and a few central documents from the cluster; and snippets--list of documents showing snippets of the document with search terms in the document highlighted in the view.

Results can be fed back into the search query to further refine the search. As with the Information Grid, this is accomplished by selecting relevant documents, then marking them to be used as a seed for future searches. Once documents are located, they can be placed in a holding area for future use.

.c4.BRAQUE. Pedersen [, 1993 #74] described the BRAQUE project, which implements a windows-based browser to a query database (built with relationship lattices) to return bibliographic information. The information is then displayed in multiple windows that attempt to give the user some idea of the underlying lattice structure of the data. The data are displayed in a hierarchical text format, with subcategories below and indented under each main category.

.c4.Generalized fisheye views. One of the earliest descriptions of a fisheye view for displaying relationships between items of information was given by Furnas [, 1986 #38]. In his paper, Furnas described a technique for focusing on items of direct interest while retaining some view of items of less interest. The most information is given about the item in the focus. Less information is given about items farther from the focus, while little information is given about items far from the focus. If the focus is shifted to another item, then the amount of information given about different items changes to reflect the new focus.

Graphical fisheye views of graphs. While Furnas [, 1986 #38] provided only text-based examples, Sarkar and Brown [, 1992 #59] and Brown, Meehan, and Sarkar [, 1993 #39] provided graphical examples of fisheye views. In their work, an information retrieval task was completed and items of information returned. The information was then displayed as a network picture, with lines connecting related items. The item with the focus would be displayed as the largest item in the display, with the size of items decreasing the farther they were from the focus. Other characteristics of relationship could be displayed by distance from the focus, so that both distance and size could indicate strength of relationship.

.c4.Fisheye view with variable zoom. Schaffer et al. [, 1996 #30] described an extension to Furnas's fisheye views for 2-D connected graphs with superimposed hierarchical clustering. Their work clusters nodes within a network, grouping individual nodes into small clusters, then grouping the small clusters into larger clusters, with the largest cluster being the full network. Nodes can be clustered in any way as long as a strict hierarchy is used throughout the clustering process.

Schaffer et al. [Schaffer, 1996 #30] used the hierarchy to produce clusters of the network that can be viewed at different levels of hierarchical detail. At each level above the network node level, clusters are represented as icons that can be opened to show the next level down. The fisheye method allows users to selectively magnify clusters while keeping the full network on the screen so that the context of the magnified cluster is preserved.

In general, users preferred the variable-zoom system over the full-zoom system. However, the fact that the full-zoom system offered more detail than the variable-zoom system (because a cluster fills the whole screen in the full-zoom system, but only a portion of the screen in the variable-zoom system), meant that for some applications, users were able to read and work with the contents of the cluster better with the full-zoom system.

.c4.Magic Lens filters. Stone, Fishkin and Bier [, 1994 #40] described a tool to view an underlying model of visible data. Like the generalized fisheye view, the data under the lens are distorted to make relationships more apparent. The data can be in a highly compressed visible format or can be completely different from the view of the data as viewed through the lens. As the lens is passed over the data, an algorithm specific to the lens is applied to the underlying data for the visible model, and the modified data are displayed in the lens portion of the interface.

Similar to a real magnifying glass, the Magic Lens can magnify a portion of a picture or text. As the lens is passed over a map, for instance, the portion of the map under the lens can be magnified to show more detail. If the visible area is text, it can be magnified to be more easily read.

In other applications, the lens might be set to show only certain words or to match certain variables in entries that the user wants to see. As the lens is passed over the data, only those entries in the data that contain the word or match the variables are made visible. All other entries remain in the background. For example, when viewing a map, the user could set the lens to highlight all schools. As the lens is passed over the map, only those buildings that are marked as schools would be highlighted in the lens.

In a search retrieval application, all the entries in the search space could be visible on the screen, as colors or points in the search space, for example. The lens could then be set to find entries based on some search criteria. As the lens is passed over the search space, only those entries matching the search criteria would be highlighted. The user could then call for more information about those specific entries in order to explore further.

The Table Lens. Rao and Card [, 1994 #41] described a viewer to be used for visualizing tabular data. As a fisheye technique, it involves distorting parts of the table to make the underlying data in the table more visible. As with other techniques described here, the Table Lens fuses graphics and text in order to increase the amount of information that can be displayed at one time in the table. Compared to a normal table, the Table Lens can display up to 100 times the information on a single screen.

In practice, a table is constructed containing a number of rows and columns. The table might be sorted on a specific column in order to highlight one facet of the data. The data are then viewed graphically so that values in each cell in a column are replaced by a line representing the value in the cell. Relationships can be discerned at a glance by looking at the chart. A specific row or column can then be expanded if the user wants to explore the data in that row or column in more detail. At this point, the text view of the data can be viewed in the row or column to see exactly what the data point says. This way, the user can explore the data, zeroing in on specific points to see exactly what underlies the visible relationships.

.c4.Value Bars. Chimera [, 1992 #42] described a tool that is most useful in displaying information about items in lists and tables that have multiple attributes. It is possible to sort the information in the lists or tables, but the list or table can be sorted on only one attribute at a time. Value bars placed next to the data can provide information about each attribute in the list or table without having to sort each one separately.

As an example, a disk's directory can be displayed with value bars. Files can have information relating to size and days since creation. The files in the display can be sorted by name, and on the side of the file display two value bars can be placed showing file size and days since creation. The value bars display cells corresponding to the values for file size and days since creation. In addition, the value bars show data for all entries in the list, even those that are not currently in view in the text view of the files sorted alphabetically. Each bar is also marked so that the cells corresponding to the visible files are shown. The user can get an immediate overview about file sizes and days of creation of the files. In addition, the user can click on a cell in the value bars to go directly to a file that has a size or a creation date of interest.

.c4.Treemaps. Johnson and Shneiderman [, 1991 #58] described a method of displaying hierarchical information in fullscreen displays called treemaps. In a later paper Johnson [, 1992 #46] described a specific implementation of treemaps called TreeViz. Treemaps are designed to display hierarchical tree information in a graphical format. The appearance of a treemap is similar to a Venn diagram. Items farther down in the tree are enclosed within walls representing items farther up the tree. An item at the bottom of the tree (a leaf node) is shown in a size representing either the size or importance of the item.

When applied to a structure such as a directory of a hard drive, the picture presented shows the sizes of the files and the relationship between the subdirectories and the files in the directory. If the treemap were applied to a hierarchical structure such as the Library of Congress Subject Headings, the relationships between subject terms could be visually displayed.

Turo [, 1994 #45] described a use of treemaps to display information about the National Basketball Association (NBA). All information for all players, teams, and the four divisions of the NBA are entered in a database. A treemap is then generated to display the information, sliced by division, team, then players. The user can then emphasize different aspects of information about the players, such as points scored, height, or years played, in order to make relationships clearer and to make certain players stand out. As each player attribute is emphasized, fisheye techniques are used to make the emphasized player block larger so that it can be seen and selected more easily.

Johnson [, 1992 #46] described TreeViz, a treemap that he applied to a hard drive containing more than 1000 files. All the files and the relationships between the files can be displayed on a single screen.

.c4.Dynamic queries. Dynamic queries on a database provide graphical results, in contrast to text search queries into a database that only return a text listing. Dynamic queries [Ahlberg, 1992 #43], embody the principles of direct manipulation and visible actions as listed earlier. Dynamic queries include the use of graphical widgets (such as sliders) to enter a query, and graphical output (such as a map) to display the results of the query.

Ahlberg et al. [, 1992 #43] described a dynamic query system based on the periodic table. The periodic table is displayed on the screen with sliders controlling search probes into the table. Values can be changed for atomic mass, atomic number, atomic radius, ionic radius, ionization energy, and electronegativity. As the sliders are adjusted, the values corresponding to the sliders' position are displayed. At the same time, elements in the table are highlighted if they satisfy the criteria set by the sliders.

Selection of values is immediate, and changes are visible immediately. The user can change the values and see the results of any change immediately. In addition, the graphical widgets permit the user to see the maximum and minimum value range affected by the widget, so that the user does not have to guess permissible values.

Current sliders usually deal only with numbers. To deal with ordered text, the text is usually placed in a scrolling list. The user then scrolls through the list and selects the desired entry from the list to use it elsewhere. Ahlberg and Shneiderman [, 1994 #44] describe the Alphaslider. The Alphaslider is used to display ordered alphanumeric data in a slider, in the same way that numeric values can be displayed. Letters are placed below the slider, with spacing between the letters indicating the number of entries under each letter. Text output is displayed above the slider. As the user moves the slider thumb, the text above the slider changes as text values are accessed through the slider. This allows text as well as numeric values to be used in dynamic query systems.

.c4.Dynamaps. Plaisant and Jain [, 1994 #47] described an implementation of dynamic queries applied specifically to maps. A dynamic query system is constructed with a map of a region as the display area. As the sliders representing search probes are manipulated, areas on the map change in response to the values that are returned. Different gradations of colors are continuously applied to the map areas to represent the values selected. The user can move the slider thumbs and instantly determine how changing the values changes the effect on the map.

.c4.Starfield displays. In [, 1992 #48] Williamson and Shneiderman introduced the Starfield display. Starfield displays join three important concepts: (a) dynamic queries, (b) tight coupling, and (c) all values in the search space visible on the screen at the same time.

Starfield displays are constructed so that all points in the search space are shown as points in the display field. By selecting values on which to search, the user adjusts the display so that more or fewer points are shown on the screen.

The Dynamic HomeFinder illustrates the concept with an application that allows a user to select a house. A base map of the Washington, D. C. area is presented on the screen. All available houses in the area are presented on the map as single points. Sliders and buttons allow the user to adjust distance from selected points on the map, number of bedrooms, cost, type of residence, and services provided. As the user adjusts the sliders or selects buttons, the map is adjusted to show only those houses that satisfy the selected criteria. As the values are adjusted, the map immediately reflects the selected values, allowing the user to explore houses on the map with minimal cost. If the user finds a point of interest, the point can be selected to obtain more information about the house.

In a later article, Ahlberg and Shneiderman [, 1994 #57] described a further application of the Starfield display, described as the FilmFinder. The FilmFinder displays all known films in the search space as points in a two-dimensional grid, with year on the x-axis and popularity (scale of 0 to 9) on the y-axis. Alphasliders are used to adjust the scale for variables such as title, actor and actress. Normal sliders adjust length of film. As the number of points on the screen is reduced by manipulation of the sliders, there are eventually few enough points that the information associated with each point can be expanded. At this level, there is enough free space on the screen to place the name of each film beside its point on the screen. If the user wants more information about a film, he or she can choose the film and have a dialog pop up with detailed information about the film. The user is allowed to choose values on this dialog to feed back into the search probe to further define the points that show on the screen.

.c4.Pathfinder Networks. Fowler et al. [, 1991 #56] incorporated some ideas from fisheye views in their information browser. Their system is specifically designed to display the network of relationships that can be formed when the searcher enters a natural language query. Once the relationship diagram is formed, the searcher can select one of the nodes and bring it to the center. Related nodes are then displayed, in ellipses, surrounding the focus. Nodes that are farther from the focus and weaker in relation to the focus will be displayed as smaller ellipses.

Articles closely matching the focus node are displayed in a separate viewer. The searcher can choose to move the focus in the network relationship, bringing different sets of articles forward, or the searcher can choose one of the article titles for an overview of the article or for retrieval. The system view can be switched so that article titles are displayed in the network relationship. Then related article titles are displayed, with network lines and the size of the article icons indicating their relationship to the focus node. Again, the focus can be shifted to other articles and new, related articles will then be brought into view.

.c4.VIBE. Korfhage [, 1991 #50] and Olsen et al. [, 1993 #55] described a two-dimensional document space called VIBE that is designed to handle a large number of reference points. Distances between reference points are calculated as ratios rather than as absolute values. Points representing documents are then projected onto a two-dimensional space. Documents that are closely related will form clusters in the document space. Documents that are strongly related to a particular reference point will be closer to that point.

Once the documents are placed in the document space, the searcher has complete freedom to reposition the reference points in the diagram. This might be necessary in order to clarify relationships between documents and between documents and different reference points. All the reference points are placed in a two-dimensional space, avoiding the problem of displaying more than three dimensions is avoided.

More information about each document in the document space can be obtained by clicking on the document icon. Once an appropriate document or cluster of documents has been selected, the full text of the documents can be retrieved.

.c4.<u>InfoCrystal</u>. Spoerri [, 1993 #52] described a visualization tool that can also be used as a visual query language to help users search for information. InfoCrystal has the following properties:

1. Users can explore an information space along several dimensions simultaneously without having to abandon their sense of overview.
2. Users can manipulate the information by creating useful abstractions.
3. Similar to a spreadsheet, users can ask "what-if" questions and observe the effects without having to change the framework of a query.
4. Users receive support in the search process because they receive dynamic visual feedback on how to proceed. They can selectively emphasize the qualitative or the quantitative information provided by the feedback to help them decide how to proceed.
5. Users can formulate queries graphically, and they have flexibility in terms of the particular methods use to retrieve the information. For example, users can seamlessly move between a Boolean and a vector-space retrieval approach, or they can easily switch from a keyword-based to a full-text retrieval approach. (p. 151)

InfoCrystal is based on the concept of Venn diagrams, except that it is easy to extend the visual picture to more than three variables. As an example, the following picture of a three-variable InfoCrystal is described:

1. Construct a Venn diagram with three circles.

2. Divide the Venn diagram into seven pieces based on the areas delineated by the intersection of the circles.

3. Convert the seven pieces into icons based on the number of variables intersecting to produce the piece. For example, the portion of the Venn diagram representing one variable only (no intersecting pieces) is represented by a circle. The portion representing the intersection of two variables is represented by a rectangle, and the portion representing the intersection of three variables is represented by a triangle. For a larger InfoCrystal, the number of sides of an icon represents the number of variables intersecting.

4. Place the icons in an enclosing shape representing the number of variables (triangle for three, square for four, etc.) in a position that indicates which variables are being represented.

Shape, proximity to variable, rank (distance from center), color, orientation, and size or brightness of icons are all used to relay information about the variables they represent. Because the icons in the InfoCrystal represent the results of various Boolean queries, the users can activate the queries by simply clicking on the icons. The user can activate the query represented by the icon, or, by choosing several icons at once, can activate a complex query as the intersection of the individual queries.

In addition to the basic features described above, users can modify the queries in the following ways:

1. Hierarchical structures can be built where multiple InfoCrystals interact. The output of one InfoCrystal can be used as the input to another InfoCrystal.

2. Retrieved documents can be used as an input criterion.

3. Users can modify the search result with the use of sliders that represent relevance weights and selection thresholds. By adjusting the sliders, the user can determine which icons are selected.

4. The InfoCrystal layout can be changed to a "Bulls-Eye" layout that places relationships with a higher relevance score closer to the center of the InfoCrystal.

5. The InfoCrystal layout can be shown in the "Bulls-Eye" layout with all documents rather than just the relationship icons.

.c4.The VisDB system. As with Starfield displays, the VisDB system [Keim, 1993 #53; Keim, 1994 #77] attempts to place all points in the database on the screen at the same time. The authors recognized that there is an upper limit to the number of points that can be displayed, relating to the number of pixels in the display portion of the screen. For search results consisting of more points than can be displayed, each pixel on the screen would represent more than one point.

To start the system, the user specifies a single query. Once the user obtains a result, the display can be modified with other query terms or by manipulating the range of the query.

The user enters a search query, returning a set of data as a result. The relevance of each item to the search term is then calculated and a value assigned to the item. A display point is then generated for each item returned and the points displayed on the screen depending on their relevance value. Those items with a greater relevance are placed nearer to the center of the screen. In addition, color is used to differentiate between groups of relevance values, with the most relevant items colored yellow, a shading to green, blue, red, then to black as least relevant. By looking at the display, the user can immediately determine the proportion of documents returned by relevance to the search term.

A specific color is related to a specific range of relevant values. These values are controlled by sliders which correspond to each of the search terms. The user can vary the range of values covered by each color, thereby modifying the search result and the corresponding display. The user can thus gain immediate feedback for additional search probes into the returned data set. Different kinds of sliders can be set up for numbers, ordinal values, nonnumeric data, or for special data types. To focus more closely on one area in the display, the user can select a specific color and enlarge its area. Once an interesting set of items is found, the user can view the specific items selected.

Current Three-Dimensional Information Retrieval Systems

.c4.MUE. The Museum Unit Editor [Travers, 1989 #60] is a front-end to the Cyc database at the Massachusetts Institute of Technology. The Cyc project has been described as follows:

> [It is] an effort to build a very large knowledge base that encompasses a broad range of common-sense knowledge. The knowledge base consists of a network of interrelated units (frames), each of which corresponds to a thing to be represented. These can be physical objects, abstract concepts, classes, or anything else of interest. (p. 147)

Information in the database is organized as an unrestricted labeled graph, but parts of it are defined so as to form directed acyclic graphs. MUE uses these to provide navigational skeletons. The directed graphs have relationships of more general to specific and more specific to general. The general-to-specific relationships can be arranged in the form of a tree in which every arrow points downwards.

The metaphor chosen to display these data was that of rooms in a museum. Every unit in the database is mapped as a room in the museum space. In order to display an individual unit, the tree must be rerooted to place the selected unit at the top of the tree.

Arrows that point up in the new arrangement can be replaced by negative arrows that point downward to form a new tree relationship. This can now be displayed as a series of enclosed spaces, with each level in the tree at the same level of enclosing room. Rooms that point from general to specific can be displayed as a different color from those that point from specific to general, to show the relationship between units of knowledge. Because the knowledge contained in the database is organized as a network, in some sense all items in the database can contain every other item in the database, so every item in the database can be displayed as a room that eventually leads to every other room.

Navigation through the database consists of following leads and relationships shown by the rooms in the "museum." To follow a lead, the user simply opens a room and moves to another level. All movement through the database is through hyperlinked navigation.

.c4.Cone Trees. Robertson et al. [, 1991 #63], described an information visualization technique for visualizing hierarchical information. The hierarchical information is displayed in 3-D in the shape of a cone and is displayed either vertically or horizontally. At the top of the cone is the item that is at the top of the hierarchical tree. Leading down from the top are items farther down in the tree. At each level down from the top, items are arranged in a circle below the item they are below in the tree. The cone shape is formed by lines connecting the parent node and the child nodes. The lines are filled in with translucent shading to make the cone shape more apparent.

In the vertical Cone Tree, the top node and its children form a shape like an umbrella. Items at the edge of the umbrella are arranged around the circle in three dimensions. Clicking on any item in the circle of items causes the cone to rotate and bring the item in focus to the front. Any item with child nodes has a further cone radiating down from it. Clicking on any item rotates the whole Cone Tree so that the item of interest is brought to the front.

In this scheme, any hierarchical relationship can be displayed at once, with the Cone Tree able to handle up to 1,000 nodes with 10 layers. The viewer can also hide cones or relationships in order to make the relationships in other nodes stand out.

.c4.The Information Visualizer. As described by Card, Robertson, and Mackinlay [, 1991 #64], this is a conglomeration of several different graphical display techniques. The immediate view is that of 3D Rooms. Each Room is a separate workspace with information relevant to a specific task enclosed in that Room. Within a Room, information display techniques such as the Perspective Wall, Cone Trees, or scientific data sculptures can be displayed, allowing further exploration of the information in a clearly prescribed context.

In order to work with the information in a specific Room, the user selects that Room with the mouse. Navigational aids are provided to allow the user to move freely from Room to Room. Color, lighting, shadow, transparency, hidden surface occlusion, continuous transformation, and motion cues provide different points of information about the data contained within a Room.

Card, Pirolli, and Mackinlay [, 1991 #65] have tested the Information Visualizer against traditional retrieval methods using a technique called the Cost-of-Knowledge Characteristic Function. This function is an attempt to work out a method to measure the information access in direct-walk interactive information visualizations. These types of systems provide for direct information access using only mouse selections and key selections.

A calendar based on the ideas embodied in the Information Visualizer was tested against a traditional UNIX calendar. In this experiment, the unit of cost being considered was primarily the user's time. The question put to the user was to view the data from a specific day, where the day to be viewed was at different lengths of time from the current day. The cost to use the Information Visualizer calendar increased arithmetically, whereas the cost to use the traditional calendar increased geometrically. Thus, the traditional calendar took a shorter time to access the day when the day to be accessed was relatively close, but the Information Visualizer calendar took less time when the day to access was far from the current day.

.c4.<u>The Perspective Wall</u>. Mackinlay, Robertson, and Card [, 1991 #62] described the Perspective Wall as one of the techniques of information visualization that can be used in the Information Visualizer. It is designed to be most useful in displaying information that can be laid out along some linear component. For example, files that have been stored on a hard disk can be displayed in the order of creation. Similarly, items that have been retrieved on the basis of one or two search terms and that can be described as a distance from the terms can be placed on a linear structure. A second dimension can be added to differentiate between the items retrieved. In the case of files on a hard disk, the second dimension could be the type of file retrieved.

The Perspective Wall is displayed as a wall with three sides joined at an angle of about 70 degrees. The middle section faces the viewer, with the other two sides angled away from the viewer in a 3-D appearance.

In the display, an item can be selected as the focus of interest. The Wall is then rotated so that the item is at the center of the screen. Other related items are then displayed on the Wall at intervals from the focus, indicating their relationship to the focal point. Only the most closely related items are displayed on the center section of the Wall. If the viewer wishes to highlight another item, selecting that item causes the Wall to slide around so that the selected item is brought to the front and center.

Relationships between items in the system can be easily inferred, with the user retaining the power to focus on any particular item.

.c4.<u>GUIDO</u>. This system, described by Korfhage [, 1991 #50], is designed to work with reference points that are numerical values assigned to documents. For example, a document could be assigned a value based on the number of times a search term occurs in the target document. The searcher defines several points to use for the search. The system then returns those documents that show a value for any of the points. If two points are used, then a flat document space is defined, with the documents placed in the document space depending on their numerical relationship to the reference points. The higher the numerical value, the closer the document to the reference points.

If three points are used, then a three-dimensional document space will be formed. N reference points will produce an n-dimensional document space. Because more than three dimensions are difficult to display, colors, shape, size, movement, and sound are required to represent more dimensions. The position of the documents in the document space shows the relationship between the documents returned.

The user can view the document space and determine the level of relationship that appears to be useful. Once a level of relationship is determined, the searcher can set a cap on the document space and obtain all the documents enclosed in the newly capped document space. Relationships are easy to determine using GUIDO as long as the number of reference points does not grow too large.

.c4.<u>Bead</u>. Chalmers and Chitson [, 1992 #54] described a three-dimensional information retrieval viewer called Bead. A large number of variables can be used to describe documents in a database--either determined manually or automatically. Latent semantic analysis and multi-dimensional scaling are then used to set up rules of physical behavior for each document. In effect, each document becomes a particle in the document space, with force and motion relationships produced between particles.

When the user enters a search request, a set of documents is returned, their positions determined by the rules calculated earlier. The documents are represented as a cloud of numbered points in a three-dimensional document space. Documents that are most closely related to the search request are highlighted in a sphere of interest, while other documents are dim. Documents tend to form clusters in the cloud of documents in the document space.

The user can zoom in on any particle in the document space in order to see neighboring particles and how they relate to each other. When a particle becomes the focus, a new sphere of interest is formed and a new set of particles is highlighted. The user can move through the document space, investigating relationships among the particles in the space.

In addition, it may be possible to rotate the document space to uncover other relationships. Particles that appear to be closely related on one dimension in the three-dimensional space, could be revealed to be far apart on other dimensions when the document space is rotated.

The searcher can select any particle, in order to investigate information about the document represented by that particle. Once a set of particles has been determined to be of interest, the documents represented by those particles can be requested by the user.

.c4.LyberWorld. Hemmje, Kunkel and Willet [, 1994 #51] designed an interface to make use of 3-D visualizations into a database of fulltext documents. It combines ideas from Cone Trees and the Vibe system and extends the Vibe system into a 3-D Relevance Sphere for display of relationships.

The contents of the documents in the system are analyzed, and the documents are placed into a hypertext relationship based on the analysis. Each document can have links to many other documents in the system, creating a tangled web of interrelationships. In use, when a search term is entered to start the search process, the web is converted to a tree in a similar manner to that employed by MUE.

When a user enters a search term, that term is made the root of the tree, and all documents and search terms directly related to that term are then transformed so that they appear as child nodes off the root. All closely related nodes are then similarly transformed and placed into the tree. The rest of the web could be similarly transformed and placed into the tree, but only a small proportion of the entire web is necessary for initial display. Once the tree is formed, the first few levels are displayed in a Cone Tree. The user can then browse the Cone Tree looking for interesting titles or interesting pathways down the Cone Tree, following a related search term down a level in the tree. As the user moves towards unexplored regions, the Cone Tree can be unfolded to show more pathways (at which point the calculations are done to transform the web into more pathways in the tree). If the user finds an unfolded path uninteresting, the path can be reduced again. If the user finds a title of interest, the document can be opened to determine if it really meets the user's needs.

After exploring some levels of the Cone Tree content space and unfolding areas that are interesting, the user can switch to the Relevance Sphere. The term nodes selected, with their related nodes, and the documents relating to the term nodes are placed into a sphere with the term nodes on the surface of the sphere and the related documents inside the sphere. Relevance relationships for the term nodes are calculated as in the Vibe system. However, whereas Vibe displayed the terms and documents in a 2-D space, the Relevance Sphere introduces a third dimension in order to clarify relationships between documents and term nodes. The more strongly related a document is to the term nodes, the closer it will be to the surface of the sphere. Three types of node/sphere manipulations can be done in order to investigate document relatedness.

1. Term nodes can be moved around on the sphere in order to see which documents move with the nodes and to investigate the interaction between term nodes.

2. The sphere can be rotated in order to separate clusters of documents.

3. The "attraction" of term nodes can be increased or decreased (as with Treemaps) to separate documents based on how strongly they are attracted to that particular term node.

At any point in the manipulation of the nodes and documents, the user can switch views from the Cone Tree, to the Relevance Sphere, to a view of the document itself. In whatever view, the context of the search is preserved and transferred into the other tool's visualization metaphor. The user is free to conduct the search using whichever tool is the most comfortable and which will yield the best results at the time.

Standards for Graphical Information Retrieval Systems

How well do the systems described here conform to the standards for graphical information retrieval systems as described by Olsen et al. [, 1993 #55]? (see The Graphical Solution section).

All the systems present their data graphically, and for reasonable amounts of data, all the data are displayed in a single screen.

Positioning of a document's icon in the display of every system gives some information about the relationship of that document to other documents in the system. In some cases, as with Cone Trees, the relationship is hierarchical. In other cases, as with Bead, InfoCrystal, or LyberWorld, the position of the icons shows how far apart the documents are based on some characteristic defined by the system.

Data reduction is used by every system. In all cases, this serves both to fit the information onto a single screen and to reduce the amount of information that the user must absorb when trying to make sense of the relationships displayed. However, all systems provide a mechanism for the user to display more information about a single item in the system. In many cases, selection of the item pops up a little window with more information about the item. In some cases, selecting the item allows the actual content to be displayed in a window on the screen. It does not appear possible to create a system that displays all important document attributes while simultaneously displaying hundreds of data points on a single screen without making the display impossible to read.

All systems provide for the entire retrieved collection to be displayed on a single screen along with the relationships between the documents except for the MUE and LyberWorld. However, in both cases, the entire collection is available for the user to jump to via hypertext links. Navigation features make the entire system available, but present only a small portion in order to make the system more manageable. As with the other systems, many more points are displayed than in conventional systems.

All systems provide the capability of highlighting a single item and retrieving and displaying more information about the item and retrieving and displaying the item itself.

All systems provide the capability to change the display interactively. Some systems, such as Cone Trees or the Perspective Wall, allow the user to grab portions of the display and interactively drag the display to highlight a different segment of the display. Other systems, such as the Starfield displays, Bead, VIBE, InfoCrystal, LyberWorld, or VisDB, allow the user to add or subtract search characteristics and to watch as the display changes accordingly. In most cases, the principles of direct manipulation are applied so that sliders or direct dragging are applied to modify the display. This allows the user to explore without having to know the exact values that are stored in the database.

All of the systems allow the user to highlight specific areas of the display. Many of the systems operate in the manner of fisheye views in that they allow the user to focus on or emphasize a single point and the system changes the display to emphasize that set of points. When this happens, the display is modified in some way to reflect the new importance placed on that particular point. In most cases, the display is not distorted, as is the case with classical fisheye views. The display can be rotated to bring the item into view (as with Cone Trees) or it can be skewed to emphasize a set of points (as with LyberWorld) or the whole visible area can move to a different area of the display space to bring forward different items (as with MUE).

Direct manipulation is a guiding principle for all of the techniques and systems described. In some cases (as with Starfield displays and the VisDB system) manipulation is accomplished via sliders that act on the visible data. In other cases, the display can be manipulated directly (as with Cone Trees, the Perspective Wall, or LyberWorld) so that the display can be rotated or changed by dragging items of interest around in the display.

All systems perform dynamic queries so that any changes made to a control are reflected in the display immediately. In some cases, the display is rotated or distorted to emphasize certain items. In other cases, the number of items visible on the display is changed to highlight one subgroup over the full display. Unfortunately, the requirement to allow novice users to begin working with little training is violated by some of the systems. Most of the current two-dimensional and three-dimensional information retrieval systems require training. InfoCrystal, in particular, will probably require extensive training.

Visual information seeking, as narrowly defined by Shneiderman and others [Kumar, 1997 #92; Ahlberg, 1992 #43; Ahlberg, 1994 #44; Ahlberg, 1994 #57; Williamson, 1992 #48], specifically calls for the use of Starfield displays for information display. However, the characteristic of Starfield displays is that they show all data points on the screen at one time. From that perspective, all the systems described here satisfy this requirement. In addition, all systems tie the manipulation of their controls to immediate action on the display.

It appears, therefore, that all the systems satisfy most of the criteria for good graphical information retrieval systems. The following section examines how well these systems address the limitations of traditional information retrieval systems, as enumerated by Korfhage [, 1991 #50] (See Introduction).

Limitations of the Information Retrieval Systems

Many of these systems still present a limited set of documents. The user enters a list of search criteria and the system returns a set of documents. The systems all display the documents in new and interesting ways, but the set of documents is still determined by the system. Even the Starfield systems, which display all the data points in the search space, are at present limited to a few thousand items.

A few systems do have the capability of allowing the user to graphically explore beyond the narrow confines of the original search. Chief among these are the fisheye system described by Fowler et al. [, 1991 #56], MUE, and LyberWorld. These systems are designed to allow the user to go beyond the confines of the visible display and extend the search space throughout the whole database. GUIDO, Bead, and InfoCrystal also have the potential of allowing the user to graphically explore beyond the current display, but it is not clear that those capabilities are built into the current systems.

All the systems display the documents in a graphical display space. The positions of the documents in the display space indicate the system's evaluation of each document on some user-defined criteria. All the systems display relationships between the documents returned, determined by the documents' positions in the display space.

The graphical displays attempt to give the user an indication of the position of a document in the corpus of documents searched. Documents closer to the center of focus are deemed most important, while documents farther from the focus are of lesser importance. However, unless the user is allowed graphically to extend the search beyond the set of returned documents, it is still difficult to tell exactly why one set of documents was selected and not another. The fisheye system of Fowler et al. [, 1991 #56] does the best job of making the criteria for retrieval clear to the user by presenting the lines of relationship used to effect a search.

Whether the full document, or only a bibliographic record, is available is still system dependent. All the systems described here have the capability of returning the full document if it is available.

It is still difficult for the user to tell whether all relevant documents have been returned. The relationships between returned documents may be clear, but not so clear is the relationship to documents not retrieved. Systems that allow the user to graphically explore further offer the most hope that the user will be able to find and retrieve other relevant documents that were not retrieved during the initial search probes. If set up correctly, Starfield displays and VisDB, with their reliance on sliders to set the parameters for retrieval, can make this process clearer. However, Starfield displays are limited to criteria set by the designers of the systems. VisDB relies on the user to set up the criteria for the sliders, so unless the user knows exactly what is in the system, it is not possible to be certain that the correct search parameters are being used.

All the systems allow the user to customize the display by changing the focus of the system. The user can change the focus to any point of interest on the display and, in most systems, can hide or show more information, depending on individual needs.

Because all the systems allow some graphical manipulation of the system, manipulating the system changes the weight of the query automatically. As mentioned earlier, however, only a few systems allow the user to graphically slide out towards and beyond the edge of the system and have the query change and extend the search space automatically as the focus moves. The user does not even need to know that a different query is being formulated, although a continuous display of search characteristics can be displayed in a separate window so that the user can monitor the search.

In all systems, the user has a large part to play in exploring the system and tuning it to display the information in a way that best suits the individual user. Again, for most systems, direct manipulation is limited to the set of points returned by a traditional search query. However, the user can usually manipulate those points in a wide variety of ways.

Although no one system adequately addresses all concerns, the graphical approach appears to offer powerful advantages over traditional text-based retrieval lists. In addition to the information provided by text-based lists, graphical systems provide the capability of displaying more information and of manipulating the information in such a way that the user's visual processing system plays a part in the information-filtering process.

None of the systems described here is the perfect information retrieval system. Even with the best of these systems, the searcher can not be sure that all relevant documents have been retrieved. The graphical systems do, however, make it easier to filter the set of returned documents to determine which documents might be relevant to the searcher's needs by making the display of relationships clearer and by making it easier to explore those relationships. In addition, the systems with the most power to influence the relevance of returned documents also appear to be those that are hardest to use. Starfield displays come closest to combining power with ease of use, but it is not clear that they can be expanded to handle large databases.

Ideal Characteristics of a Graphical Information Retrieval System

Each of the systems described has strengths and weaknesses when considered as a graphical front-end to an information retrieval system. Current research, indicates that the following characteristics are among those that should be incorporated into a graphical information retrieval system:

1. The system should be visual and graphical.

2. Documents in the system should be represented by icons that can be different sizes, shapes, colors, and/or have sounds associated with them.

3. The system should be based on the fisheye principle. When the user focuses on a single point, other points in the system should dim in proportion to their relationship to the focus point.

4. The display of information should potentially be three-dimensional. The third dimension gives the capability of displaying relationships that cannot be seen in a simple flat display. This recommendation is tempered with the realization that adding a third dimension adds to the complexity of the system.

5. The system should allow the user to move and rotate the display on the screen, bringing different parts of the display to the front on command. It should be possible to highlight several different parts at once to explore the relationships between disparate parts of the display.

6. The system should allow the user to explore beyond the confines of the original search by bringing parts of the display at the periphery to the center and thereby extending the search into previously unknown regions as new pieces of the document space are brought into the display.

7. The system should allow the user to view relationships between several different kinds of objects. At a minimum, the user should be able to change the display to display relationships between search terms or to display relationships between actual documents.

8. The system should allow the user to select and retrieve any number of documents by simply pointing to the document in the search space and issuing a simple command.

9. The system should be composed of several parts, consisting of the relationship display space, a partial list of returned documents, the user's original query, a display showing the initial system-derived search request, and a display showing the current system derived search request.

10. The system should allow the user several different ways to look at the same data set.

Would an information retrieval system based on all these characteristics be the ideal system? Possibly not, but the power that current graphical systems give to the individual user to control the direction and extent of the search probably means that they come closer than any nongraphical system in use today.

Virtual Reality

Virtual reality (VR) systems offer the capability of placing the user inside the search interface itself. In comparison to the conventional graphical search interfaces described here, VR systems should be able to offer the same capabilities, with the added immersion factor. The following section describes the capabilities of VR systems and the use of Virtual Reality Modeling Language (VRML) to add VR capabilities to the World Wide Web.

> What is VR good for? Flohr [, 1996 #13] presented the following scenario: Imagine a travel agency's home page. Instead of a boring list of choices or photos of destinations, you find yourself in a 3-D scene – a virtual travel agency. While waiting for the "agent," you step up to a poster of a resort. A mouse-click takes you into a scene in the poster, complete with trees, mountains, and a country club. Ordinary, flat Web pages will seem snore-inducing by comparison. In addition, the "no-interface" interface will appeal to the neophyte net surfer. (p. 61)

<u>Defining Virtual Reality</u>

.c4.<u>Conventional definitions</u>. There is no commonly accepted definition of VR; in fact, the definition of VR depends very much on who is asked. For those working on high-end systems, VR is synonymous with immersion in the computer environment. Instead of looking at a scene, "the person who experiences VR is surrounded by a three-dimensional computer-generated representation, and is able to move around in the virtual world and see it from different angles, to reach into it, grab it, and reshape it" [Rheingold, 1991 #1, p. 17]. Similarly, Pimental and Texeira [, 1993 #2] defined VR as "a human-computer interface in which the computer creates a sensory-immersing environment that interactively responds to and is controlled by the behavior of the user" (pp. 14-15). For Pimental and Texeira, the key words to describe VR are <u>sensory-immersing</u> and <u>interactive</u>. Not only must the user be able to interact with the virtual environment, but the user must be immersed in the environment as well.

Stuart [, 1996 #69] used virtual environment instead of VR to differentiate between high-end virtual environments versus low-end VR systems. Ellis [, 1994 #23] defined virtual environment displays as "interactive, head-referenced computer displays that give users the illusion of displacement to another location" (p. 17). In virtual environments, the user is immersed in a VR environment that provides an interactive and multisensory experience. The environment uses high-quality (usually expensive) devices to provide real-time updates, and position sensing to provide a realistic simulation of the environment. The key is that the user is immersed in the environment and interacts only with objects in the environment.

.c4.Reasonable definitions. The problem with conventional definitions of VR is that they focus on the immersive experience. This means that the user must don head-mounted displays and body suits or be surrounded by a whole room full of equipment to experience conventional VR. However, as Hand [, 1994 #71] pointed out, full-immersion VR systems cannot reproduce the physical environment in any way close to reality because of the constraints of computing power and sensation-producing technology. He thinks that the pursuit of VR systems that would be indistinguishable from physical reality may be the wrong approach.

Hand [, 1994 #71] preferred to define VR as "something which is not real, but may be considered to be real while using it" (p. 107). This allows the system to contain elements that give the user the illusion of reality, without having to provide the full fidelity of physical reality. For example, books, plays, and movies do not replace physical reality, but participants in these forms of entertainment still have some sense that the experience is real.

Latta and Oberg [, 1994 #80] agreed with this assessment.

> We feel that the current fixation on immersive VR systems fails to encompass the potential of the technology. The core of VR is the human-computer interface. A VR system might range from a flight simulator to a synthetic environment to

telepresence. The definition and conceptual model for the technology should be robust enough to encompass all three forms, as well as others that develop as the technology evolves. (p. 23)

For Latta and Oberg [, 1994 #80], immersive VR is just one aspect of total VR. In fact, as long as VR systems do not fully replicate human senses, “then the experienced model will be a combination of the artificial and the real environments” (p. 28).

For Machover and Tice [, 1994 #81], the core of VR is the emphasis on the experience for the human participant. It is simply a new way to look at human-computer interaction.

Stuart [, 1996 #69] listed classes of VR systems that do not meet the definition of conventional VR:

CAVE systems--the user is placed in a room with VR images on all the walls.

augmented reality--sensory stimuli are superimposed over real world sensations.

vehicular simulators--vehicle cockpits have views projected on the windows of the simulator.

fish tank VR--VR images are displayed on a conventional monitor.

Uses of Virtual Reality

.c4.Current uses. Because immersive VR systems are usually expensive, “most virtual reality applications have been built by or for large organizations’ specific internal applications, and even these are more often than not proof-of-concept prototypes, rather than something that is in everyday use” [Deering, 1995 #34, p. 221]. A few successful applications have appeared, such as virtual architectural walkthroughs by Matsushita Electric Works of Osaka and VR video games from Virtuality Entertainment Systems [Ellis, 1994 #23].

Beyond the few corporate VR systems, there are some academic systems. For example, Deering [, 1995 #34] described the HoloSketch system, which uses head-tracking stereo glasses and a 6-D hand input device to control the manipulation of objects in the 3-D system to draw in three dimensions. Dede, Salzman, and Loftin [, 1994 #73] described Newton World, a VR environment designed to teach Newtonian physics.

.c4.Potential uses. Despite the fact that there are very few existing VR systems built to date, most authors are confident that there are many potential uses of VR. For example, some uses for VR might include virtual travel, 3-D games, telepresence, telesurgery, positioning radiation beams for treatment, architectural design, games, virtual theater, virtual art museum, virtual classroom, 3-D subject lessons, dangerous lab experiments, armchair tourism, flight simulation, battle planning, erotica, study of insanity, and fantasy [Larijani, 1994 #68]; outdoor landscape design, artistic designs, business money flow, stock market data analysis, chemical and molecular simulation, virtual prototyping, virtual bodies, and auto test driving [Pimental, 1993 #2]; and entertainment, virtual humans, molecular modeling, concept design, kitchen design, and virtual manufacturing [Harrison, 1996 #70].

Ellis [, 1994 #23] also feels that "virtual environment displays potentially provide a new communication medium for human-computer interaction. In some cases, they might prove cheaper, more convenient, and more efficient than former interface technologies" (p. 17).

Virtual Reality in Practice

.c4.Problems with virtual reality interfaces. VR has found some practical demonstrations, but at a cost that has placed it well beyond the reach of most in the scientific and academic community. However, with the increase in the power of high-end desktop machines, VR experiments are now within the reach of most of the scientific community [Wann, 1996 #28].

However, there are problems with today's desktop VR applications. According to Waterworth and Serra [, 1994 #24]:

> Most of today's user interfaces for standard 3-D applications do not make much use of the third dimension. The result is that users interact with variations of

> standard, 2-D interface objects such as menus. No attempt is made to capitalize on the possibilities for enhanced HCI that the extra dimension provides, partly due to the absence of appropriate design models. (p. 319)

In addition, current virtual environments can produce nausea and altered visual and visuomotor coordination with extended use. These problems can cause interference with normal activities [Ellis, 1994 #23]. "Current virtual reality technologies have not yet crossed the threshold of usability. Display resolutions in may cases render the user legally blind. Head- and hand-tracking devices are inaccurate and of very limited range" [Ribarsky, 1994 #82, p. 10].

Simulator sickness and side effects are common after one spends a lot of time in a VR environment, with rates of some kinds of side effect ranging up to 90%, depending on the simulator. Usually the side effects include some kind of nausea, but sometimes dramatic disorientation can occur, to the extent that people lose control of fine motor skills or suffer flashbacks to the simulator environment [Strauss, 1995 #27].

There are also problems readjusting to the real world for users who remain in the simulation environment for too long. As Kreuger [, 1996 #26] noted, "Slthough the severity of simulator sickness declines with experience in the simulator, readjusting to the real world upon leaving the simulator -- what I term <u>reality sickness</u> -- becomes more of a problem as time goes by" (pp. 1-2).

Another problem is that the motions necessary to operate in a virtual environment for example gesturing, walking, and pointing, can be more physically tiring than the motions needed to control a one- or two-dimensional interface. VPL (a company that made data gloves for use in VR environments) went into bankruptcy because the VPL DataGlove failed to sell as well as the company expected. The Mattel PowerGlove, a derivative of the DataGlove, failed because its applications proved physically tiring [Ellis, 1994 #23].

Unfortunately, "virtual reality has been the focus of a large number of media reports that have portrayed the more speculative aspects of the technology" [Wann, 1996 #28, p. 829] without considering the problems that users experience in the virtual environment.

.c4.Problems with the development of VR interfaces. Enthusiasts talk as if immersive VR will be the environment of choice for general computing. The reality is quite different. "Stephen Ellis claimed last year that in the past 30 years the only major commercial product to emerge from virtual environment or teleoperation technology is flight simulation" [Bolter, 1995 #79, p. 8].

Deering [, 1995 #34] observed the following:

> Virtual reality (VR) is often promoted as the next generation of man-machine interface, but other than in simulation and entertainment, virtual reality has not yet found widespread use in everyday applications. Any application in which rapid, accurate understanding of the shape, orientation, and distance to 3D [sic] objects is important is a candidate for productivity improvements through virtual reality. But there are many barriers to deploying virtual reality technology. Only recently has 3D [sic] graphics hardware of minimally acceptable rendering performance become affordable. Trackers and displays are fraught with technical limitations. Building any completely new software application has become progressively more expensive in recent years, as user's expectations of functionality, quality, documentation, and support have steadily risen. Standardization of any form, whether hardware or software, is not yet on the horizon. (pp. 220-221)

One of the claims for VR is that it should be just like the real world, so that, in effect, the interface will effectively disappear. "This is neither possible in the short run nor desirable in the long term" [Krueger, 1996 #26, p. 1]. According to Bolter [, 1995 #79], "Unmediated interaction clearly does not suffice when we introduce symbolic information, such as texts, numbers, graphics, or even graphics used symbolically (as in scientific visualization)" (p. 9) The interface will not disappear in the near future, so we must make sure that whatever we design will be worth using. As Kreuger [, 1996 #26] says, "many of the ingredients of that [encumbering and intrusive human VR] interface have indeed disappeared from use because they are not yet worth the inconvenience of wearing them" (p. 1).

Dede et al. [, 1994 #73] listed what they consider to be some of the challenges to the development of good VR systems:

> Virtual reality's physical interface is cumbersome.
>
> Display resolution is inversely proportional to the field of view. The low resolution of current VR displays limits the fidelity of the synthetic environment and prevents virtual controls from being clearly labeled.
>
> VR systems have limited tracking ability with delayed responses.
>
> Users often feel lost in VR environments. (p. 90)

Unfortunately, the press does not seem to understand the limitations of current technology and coverage has led many potential users to overestimate the actual capabilities of existing systems. VR developers must work much harder to overcome the limitations of current systems and to develop systems that will actually meet the needs of users. However, "unless their expertise includes knowledge of the human-machine interface requirements for their application, their resulting product will rarely get beyond a conceptual demo that lacks practical utility." [Ellis, 1994 #23, p. 18]

Another aspect of the press coverage has been the encouragement of interest among nontechnical groups and organizations. Some of these groups have even sponsored conferences and workshops, but the level of expertise and the content of the conferences has often varied widely. These proponents of VR have claimed that VR is “a very special field where there are no experts, and everyone can be one” [Ellis, 1994 #23, p. 18]. Ellis [, 1994 #23] noted that nothing could be more false. The research and development community associated with practical VR development efforts have the technical training and applications background required to design usable virtual environment displays and to constitute a tradition of expertise in this field.

For traditional VR systems, there is also a requirement for a significant investment in hardware. There are, in fact, relatively cheap systems, but this cost is relative. Ellis’s [, 1994 #23] comment that “working virtual environment display systems in the moderate to low price range (that is, under $60,000) remain isolated to date” (p. 21) indicates the level of commitment that even the low-priced hardware requires. There are even more problems with cheaper hardware than with high-priced systems. According to Ellis [, 1994 #23], “Most systems using the cheaper accessible technology have failed to pass beyond the stage of conceptual demonstration to the stage of useful work” (p. 20) The ability of the cheaper technology to simulate reality is just not good enough to do good immersive VR.

Even with expensive systems, the gap between what a computer system can simulate and the amount of physical perceptions we experience in the real world is so large that, for any current system, “it will be necessary for the system to fake some of the causality and for the participant to cooperate with the illusion to make it work” [Krueger, 1996 #26, p. 1]. While it is true that high-level simulations work best, low-level simulations can also give good results, because the participant must willingly cooperate to make VR systems work, at any level of simulation. For example, Krueger [, 1996 #26] noted that while vehicle simulations provide a high-level of realism and achieve a high level of performance, other applications provide only a minimal level of reality yet still manage to achieve good results. The ability of an application to affect the user is influenced by how much the user willingly participates in the illusion behind the simulation.

.c4.<u>Development of reasonable VR</u>. In order to simulate a VR environment totally, all sensory inputs must be considered and all perceptions must be artificially created. However, this is not necessary in most cases, because it is only necessary is to simulate the environment enough to accomplish the task at hand [Wann, 1996 #28].

Because the participant is willingly cooperating in the illusion of VR, all that is necessary is to create an experience that the user can make real, rather than to create an experience that is indistinguishable from reality. "In this sense, novels, theater, and film already provide passive experiences that the readers or viewers feel that they have lived through" [Krueger, 1996 #26, p. 2]. Krueger alo said that most of what we experience in modern life has already been modified from the original physical reality.

> Although we tend to think of physical reality as natural, most of what we encounter in our modern lives is not. We have constructed physical reality to accommodate the lives we want to live and the functions we want to perform – to the point where patterns of our activity are visible in satellite images. [Krueger, 1996 #26, p. 2]

According to Latta [, 1994 #80] , "Realism is a qualitative judgment, achieved when the participant perceives equivalence between interactions with an artificially created environment and interactions with a real environment" (p. 24) As long as the user is participating in the illusion, “the experience does not necessarily have to be realistic--just consistent” [Machover, 1994 #81, p. 15].

For most people, using VR to simulate such things as a supermarket with virtual aisles and virtual food is counterproductive and pointless. Simple lists work much better and are much faster. The goal is to build environments that require three-dimensional characteristics to really evaluate the data [Wann, 1996 #28]. For example, structures that are three-dimensional in the physical world can be ambiguous when projected onto a 2-D computer monitor. VR techniques that include the ability to rotate the objects and to view the objects from different angles make it easier to remove ambiguity from the 3-D structure [Encarnacao, 1994 #37]. “The compelling aspect of VR is the experience it creates” [Latta, 1994 #80, p. 24]; the ability to allow the user into the simulation is the key.

Once the user is placed within the three-dimensional space, the capability of moving through that space is a requirement. “A two-dimensional flow chart can be adequately viewed from a single viewpoint or by scrolling the display horizontally or vertically. Levels of detail, however, are hidden within a three-dimensional structure unless the viewpoint can be changed interactively” [Wann, 1996 #28, p. 830].

There is still the problem of countering the effects of cyber-sickness within three-dimensional environments. The sickness may be due to a sensory conflict between the simulation and the experience of reality. This is especially true in situations in which the user is immersed in the VR environment. Suggestions for combating cyber-sickness include (a) reducing the immersive nature of the simulation, (b) limiting the time in the environment, (c) limit the amount of unexpected changes in motion within the VR environment, and (d) increasing the power of computers to make the simulation better [Strauss, 1995 #27]. Computing power is increasing, but even the fastest computers available today still have problems processing data fast enough for accurate simulations. It is also difficult to limit the number of unexpected changes in the environment without seriously limiting the simulation. It would be difficult to predict what each user would consider an unexpected change. Similarly, unless the researcher has strict control over the simulation and the participant's use of the simulation, it will be difficult to limit the amount of time in the simulation. This leaves reducing the immersive nature of the simulation. Because this usually involves less expense and control than in traditional immersive VR simulations, this is probably the easiest option to adopt.

It is also not clear that immersive, head-mounted virtual environments perform better than non-immersive, non-head-mounted environments. As Ellis [, 1994 #23] points out:

> A key element frequently missing in the research for many applications areas is a rigorous comparison of user performance with a head-mounted virtual environment display versus a well-designed panel-mounted substitute. Panel-mounted formats are publicly viewable, available with high resolution, and currently much cheaper than head-mounted virtual environment systems. Without such comparisons, the specific benefits of the new technology will remain unknown, and the market will wait on developments. (pp. 20-21)

In current research, Fish Tank Virtual Reality simulations come closest to providing the experience of VR while delivering simulations in a relatively low-cost and non-immersive situation. Although these simulations are usually head-referenced, they also come the closest to providing the kind of experience with VR that VRML on the World Wide Web provides.

Fish Tank Virtual Reality

Fish Tank Virtual Reality is a technique for viewing three-dimensional scenes on a conventional computer monitor. Typically, such a system uses special glasses to generate 3-D images that appear to float above the surface of the monitor. In addition, special equipment is attached to the head to allow the system to track head movements, making the system head-referenced. For example, Ware et al.[, 1993 #67] used a system that consists of a Silicon Graphics 4D 240 VGX workstation and display monitor, with StereoGraphics CristalEyes stereo glasses, and a head-mounted, head-coupled system to track head movement. This particular system is still relatively intrusive, but it does not require massive processing power to generate a relatively detailed and fast-moving scene. Although the complete system described by Ware et al. contains hardware elements that generate floating 3-D scenes and that allow for head-referenced motion, these systems do not have to contain all the elements to be considered Fish Tank VR.

Ware et al. [, 1993 #67] described a test with their system in which they looked at the effects of stereo (using the CristalEyes stereo glasses), the head-mounted, head-referenced system, and a combination of the two. They compared the effects of stereo, head-coupling, and both stereo and head-coupling while (a) viewing a 3-D scene and (b) performing a task involving tracing a path through a tree display. Results of their test showed that head-coupling was more effective than the stereo for both tasks, with stereo only marginally better than no stereo. The results suggest that head-referenced motion within the 3-D scene is more important than stereo. However, the authors also indicated that the effect may be due more to the motion of the image within the display than to actual head-coupling itself.

In a separate test of the Fish Tank VR system, Ware and Franck [, 1996 #83] looked at using the Fish Tank VR system for more abstract information, such as database schemas, networks of human relationships, the structure of object-oriented code, or hypertext links. The authors noted that "it is by no means clear that 3-D diagrams offer any advantages over 2-D ones in representing this kind of information. The present paper provides some strong evidence that 3-D diagrams can offer significant advantages under the right viewing conditions" [Ware, 1996 #83, p. 123]. The authors' contention is that, given the right kind of task and interface, Fish Tank VR can be very effective. Because Fish Tank VR is the type of VR interface closest to VRML, these findings indicate that VRML might also be effective representing abstract information in 3-D.

VRML

Introduction

> The idea of a VR interface to the WWW was born and first discussed in Geneva, Switzerland, Spring 1994. VRML is intended by its designers to be the standard language for exchanging virtual worlds (3-D scenes) including interaction possibilities and multi-user abilities via the Internet. [Broll, 1996 #6, p. 427]

Broll and Koop also noted that VRML became the established 3-D standard soon after the first Internet graphical browsers were released. VRML is a platform independent standard that requires only a browser to view 3-D scenes. As such, it could "make navigating through on-line museums, libraries, and marketplaces on the Internet as common as interacting with textual information is on the WWW (World Wide Web) today" [Vacca, 1995 #14, p. 28].

Defining VRML

.c4.What is VRML? VRML is similar to HTML in that it is a description, contained in an ASCII text file, that tells the browser how to display pages. Unlike HTML, it tells the browser how to display 3-D worlds rather than just document pages.

However, like HTML, VRML operates within the basic Web browser-server paradigm. VRML files are stored on a Web server; servers then handle browser requests and return the requested VRML document, preceded by the document's Multipurpose Internet Mail extensions (MIME) tag to tell the browser exactly what kind of file it is. The MIME type for VRML documents is x-world/x-vrml. VRML files can then be viewed with a VRML viewer and a browser configured to use that viewer [Flohr, 1996 #13, p. 62].

VRML's creators (Mark Pesce, Gavin Bell, and Tony Parisi) designed VRML 1.0 to be platform-independent, suited for low-bandwidth network connections, and easily extensible. In order to achieve this, they sacrificed other features, including interactivity. According to Flohr [, 1996 #13], "You can see that virtual basketball and move around it, but can't dribble it–-not yet anyway" (p. 61), although interactivity is promised for the next version. In addition, the designers were concerned that the files would be transmitted over the Web, possibly through a modem, so they made sure that the files also had a compact file format, allowing them to be transmitted quickly over the Web [Amdur, 1995 #22, p. 49]

VRML is a language that specifies dimensions, textures, light sources, object materials, and other instructions for creating a 3-D scene. The server transmits a file with an object command, such as "cone," and the mathematical description of its shape and location to the client, along with graphical elements such as texture and color. The client then handles rendering and manipulation of the 3-D scene [Amdur, 1995 #22, p. 49].

.c4.<u>Resources required for VRML</u>. VRML requires a powerful CPU. Animated 3-D scenes are computationally intense computing applications, even if the scene is not producing photo-realistic rendering. Without sufficient computing power, one can expect to see sluggish navigation, stuttering frames, and artifacts (such as jagged edges). VR and 3-D rendering work best on high-end graphics workstations, such as those produced by Silicon Graphics; however, for most VRML applications, a 100-MHz Pentium machine with 32 MB of memory and Windows acceleration hardware or a Power Mac will do [Flohr, 1996 #13, p. 62].

Flohr [, 1996 #13] also claims that a fast network connection is a necessity. Even though VRML describes 3-D objects efficiently, a complex scene has many polygons, and textures require much data. Flohr suggested that ISDN is the minimum to explore the 3-D Web seriously. Amdur [, 1995 #22], however, feels that the file format is sufficiently compact, that "a modem faster than 14.4 Kbps does not markedly improve performance except for the initial download of the player. But a powerful CPU does make a difference, especially if it's coupled with a 3-D graphics accelerator board" (p. 49). Even with complex files, most of the time spent in showing and navigating a 3-D the scene is done after the file is received, not during download.

<u>Uses of VRML</u>

.c4.<u>Current uses</u>. There are few VRML applications deployed today, although some architecture and industrial-design firms are using VRML to construct static models to display to customers and the general public. In addition, some companies are experimenting with using 3-D to aid in navigation around their Web sites [Flohr, 1996 #13] .

Some experimentation is being done at the Fraunhofer Institute, where VRML is used and tested as an output and visualization format for 3-D data in different application areas. The attraction of VRML is the same as that for conventional VR--it provides the potential to more clearly display data in a 3-D format than in a 2-D presentation. The hope is that VRML will make interaction with the data more intuitive for the user and also provide different avenues for exploring the data [Broll, 1996 #6].

.c4.Potential uses. Chinnock [Chinnock, 1996 #21] discussed the future use of VR technologies:

> Thanks to new development tools and hardware, digital convergence, and the Internet, vendors and analysts predict that VR technologies will become commonplace of the next two years. Three-dimensional worlds on the Internet will be popular because it's a metaphor that everybody can relate to. This is exactly what is needed to bring non-computer-literate mass markets on-line. (pp. 26-27)

According to Broll and Coop [, 1996 #6], "Using a VRML viewer, it is possible to navigate freely through the scene and to follow hyperlinks to new 3-D worlds, HTML documents or other valid MIME types by selecting linked objects within the 3-D scene" (p. 427). In addition, all the uses that are predicted for conventional VR are predicted for VRML.

Problems with VRML Interfaces

Because VRML on the Web does not involve total immersion in the environment, the effects of cyber-sickness are not a problem when interacting with VRML applications. However, VRML 1.0 is limited in function. "Currently, VRML only allows browsing through static 3-D worlds without interactive behavior" [Broll, 1996 #6, p. 433]. Also, as with the interaction with conventional VR simulations, there is a learning curve involved with navigating in the VRML world. As Flohr [, 1996 #13] has written, "Is it true then that VRML brings us the Web without the interface? Not exactly. Walking or flying through a scene may be more natural than scrolling through a Web page, but you still have to learn how" (p. 62).

Information Retrieval Strategies

Information retrieval strategies fall into two general categories: (a) searching and (b) browsing [Lin, 1997 #84 ; Thiel, 1996 #85; Rada, 1995 #88]. Searching involves the formulation of a query by the user, usually a Boolean logic expression or proximity matching expression. The retrieval system then takes this expression, matches it against a specific set of system-defined criteria, and returns a set of documents that it determines are a match for the original user-formulated query. In this case, the objective of the system designers is to generate only that set of documents that exactly match the user's intentions.

Two major problems with searching are that (a) it is often difficult for a user to form a good query that allows the system to return a good set of documents and (b) documents are usually returned in a linear form that is not easy for the searcher to deal with [Lin, 1997 #84]. The system often returns either too few or too many documents to be of worth to the searcher. Search queries are system specific, and if the searcher does not understand the structure of the system, it is impossible consistently to form good queries. To use search systems well requires training and experience. In addition, searchers often do not have a clear idea of what they are trying to find. Even when the searcher can form a query, the system returns only those documents that match its designers' ideas of relevance, not those of the searcher [Lin, 1997 #84; Chang, 1993 #86].

Browsing tries to address the two problems of searching by (a) starting with or returning a large set of documents that does not unduly constrain the document space and (b) visually displaying the documents and relationships among these documents to allow the user to choose those documents of most interest [Lin, 1997 #84].

Browsing is a useful information retrieval technique when (a) people have no defined search strategy, (b) people have a poor understanding of the information organization in the system, (c) the information system is organized in such a way as to encourage browsing, (d) information is easier to find than describe, and (e) people want to skim through the information [Marchionini, 1989 #90; Marchionini, 1987 #87; Thompson, 1989 #89].

Hierarchically organized data is a good candidate for browsing strategies. Information is already organized in such a way that all the information in the structure is available from the top node, even though it may not all be visible at once. In addition, the user can move through the tree structure, using an already defined structure to help define the query. In addition, each choice in the tree structure reduces the amount of information that the user must deal with, allowing further refinement of the search by the user rather than the system. Finally, the structure allows the user to either stop at a category or a single document and return all the information at either the category or document level [Kumar, 1997 #92; Marchionini, 1995 #91].

Virtual Reality Search Interfaces

Darken and Sibert [, 1996 #29] described some of the problems with navigating within a large virtual world. "In fact, problems associated with wayfinding and navigation are predictably encountered in every large virtual world. Navigators of these worlds often become disoriented and are unable to perform the simplest of searching tasks" (p. 50). In order to increase the user's ability to navigate successfully within a virtual world, the developer must be careful to provide sufficient clues as to the structure of the virtual world. To increase navigability of the virtual environment, Darken and Sibert [, 1996 #29] suggested that the developer provide the following:

1. Landmark knowledge--information about the visual details of the environment. It is memory that a user retains about notable features in the environment, such as mountains or buildings.

2. Procedural knowledge (or route knowledge)--information necessary to allow the user to follow a particular route.

3. Survey knowledge--configurable or topological information.

Essentially, the developer of a virtual environment is building a 3-D map of some underlying structure. A physical map is congruent to locations in the physical world. The 3-D map could also be congruent to the physical world; however, it could also be related to some underlying data structure. The closer the 3-D environment represents the user's model of the structure, the more quickly the user of the environment can learn how to navigate within the environment.

Darken and Sibert [, 1996 #29] suggested the following organizational principles to help the user mentally organize the virtual environment in order to navigate within the environment:

> 1. Divide the large world into distinct small parts, preserving a sense of place. This division should be hierarchical in nature.
>
> 2. Organize the small parts under a simple organizational principle, such as a grid or logical spatial ordering (e.g., a street-naming convention).
>
> 3. Provide frequent directional cues. This information can be supplied in the form of directional landmarks or independently, as in a compass. (p. 51)

The form of the task is also an issue. While navigating within the environment, the user has some goal in mind. If only primary, task-related, purposeful movement is considered, Darken and Sibert [, 1996 #29] list two broad classifications of possible tasks: searching and exploring, which they then extend to:

Naive search: Any searching task in which the navigator has no a priori knowledge of the whereabouts of the target in question. A naive search implies that an exhaustive search is to be performed.

Primed search: Any searching task in which the navigator knows the location of the target. The search is nonexhaustive.

Exploration: Any wayfinding task in which there is no target. (p. 54)

The 3-D interface designer must consider these three kind of searches when building a 3-D search application.

Research in Virtual Reality

There is growing evidence that representing diagrams in 3-D can allow more complex information to be comprehended. An influential work has been the SemNet project. This used a 3-D representation to allow users to visualize large knowledge bases as nodes and arcs in a three-dimensional space. [Ware, 1996 #83, p. 125]

Ribarsky et al. [, 1994 #82] report on preliminary research to display the full text of a book in 3-D space. "The purpose of any 3-D layout is to use viewers' innate sense of space to help them understand the structure and browse sections of the text. We could visualize hypertexts as well as conventional books in 3-D" (p. 12).

Jacob, Sibert, McFarlane, and Mullen [, 1994 #36] reported on the results of an experiment to test two kinds of 3-D tasks. The first task was to move a square into a reference square while simultaneously changing the size of the first square to fit the reference square. The second task was to move a square with a colored dot into a reference square, then change the color of the dot to another color. A 3-D input device and a mouse were used for each task, and the time to perform the tasks was tested. The 3-

D input device was better for the first task, whereas the mouse was better for the second task.

Humans use many different cues to estimate depth in scenes. According to Kjelldahl and Prime [Kjelldahl, 1995 #33], “Static cues include occultation, size, perspective, and shadows” (p. 199). Objects farther away may also appear darker, perhaps because of some effect such as fog. In addition, movement can help to determine depth, either the movement of two objects in relation to each other, or movement towards or away from the viewer.

Kjelldahl and Prime [, 1995 #33] reported on the results of an experiment to determine the importance of (a) complexity of rendering of the image (e.g. wireframe or filled in), (b) lighting and shading, (c) horizontal vs. vertical placement, and (d) type of object (e.g. cube or dish) on subjects’ ability to estimate the apparent depth of objects on the screen. Results showed that placement and lighting of objects were important to depth estimation, with horizontal placement providing more depth cues than vertical placement of objects. The type of object and the complexity of rendering of the images did not significantly affect estimates of depth.

Gallimore and Brown [, 1993 #35] conducted an experiment in which subjects compared the shape of two 3-D Computer-Aided Design (CAD) objects to determine whether they were the same or different. Subjects were tested with (a) simple vs. complex shapes; (b) wire-frame, shaded, texture, or shade plus texture; and (c) stereo, provided through the use of stereo glasses vs. nonstereo display only. Results indicated that simple objects were easier to work with than complex objects. Wire-frame objects were more difficult to work with than other shadings. Providing stereo cues did not result in any improvements versus nonstereo.

Research - Advantages of Virtual Reality

The studies reported in the previous section assume that 3-D interfaces are better than 2-D interfaces. However, as Ware and Franck [, 1996 #83] pointed out, most studies have not answered the most fundamental question: “How much is gained by moving from a 2-D to a 3-D representation?” (p. 125).

Ware and Franck [, 1996 #83] reported on the results of an experiment to compare a 2-D interface with a 3-D interface to the same data space. The experiment used 75 nodes with 100 arcs between the nodes. The task was to find an arc with a certain path length between two nodes. Different conditions were tested including the following:

1. 2-D--no stereo, no rotation
2. Static perspective--no stereo, no rotation
3. Stereo--no rotation
4. Passive rotation--no stereo, the scene rotated at a constant velocity
5. Stereo, passive rotation--stereo with constant rotation
6. Hand coupled--no stereo, movement was linked to mouse movements
7. Stereo, hand coupled--stereo, motion was linked to mouse movements
8. Head-coupled perspective--no stereo, movement was linked to head movement
9. Stereo, head-coupled perspective--stereo, movement was linked to head movement.

Results from the experiment showed that, in all cases, the 3-D condition produced significantly fewer errors than the 2-D conditions. Stereo plus motion produced significantly fewer errors than conditions without motion. The response times were relatively uniform across conditions. The hand and head motions were significantly slower than all other conditions, including the 2-D conditions.

Ware et al. [, 1993 #67] and Ware and Franck [, 1996 #83] reported on a number of studies that show that the number of errors in detecting paths through tree structures is substantially reduced if a 3-D display method is used. Both scene rotation and stereopsis seem to help reduce errors in path-tracing tasks. Studies have also shown that motion is more valuable than stereopsis in reducing errors when subjects were required to discover to which of two tree roots a highlighted leaf node was attached. Ware and Franck also found that motion helped to reduce errors, although in the experiments performed by Ware and Franck, the motion was caused by head movement, rather than movement within the environment itself. Again, the key to reducing error in the 3-D environment seems to be movement within the 3-D scene.

Why Research in Virtual Reality and VRML Is Necessary

> The increasing prevalence of 3D [sic] visualization application environments has also increased the necessity and difficulty of interactive exploration. These environments make interactive exploration more important because the phenomena are more complex. They make interactive exploration more difficult because they offer control of more degrees of freedom. [Encarnacao, 1994 #37, p. 68]

Wann and Mon-Williams [, 1996 #28] made the following observation:

> The question of "who really needs VR?" is not dissimilar to the question of "will everyone really use a desktop computer?" that echoed through many finance sub-committees in the 1980s. Prior to 1980, it was difficult to envisage how widespread computer use might be, particularly in occupational areas that did not hinge on "computation". Similarly, at present it is difficult to conceive of how interactive three-dimensional visualization might be used extensively in, for instance, the social sciences. (p. 830)

If VR and VRML do follow the model of the desktop computer and HTML, then VRML applications will become commonplace in the very near future.

A hierarchically structured database of information is ideally suited for experiments with 3-D information spaces. It is already structured in such a way that displaying this structure in 3-D is relatively easy. In addition, spatial data are designed to be browsed rather than searched, and a hierarchically structured system provides a good environment for this sort of task.

CHAPTER 3

METHODOLOGY

Participants

The participants for the study consisted of undergraduate and graduate students who currently use World Wide Web search engines, such as Yahoo, to search for information on the Internet, and who use a Web browser with a graphical user interface, such as Netscape Navigator. The sample for this study was drawn from students enrolled in Computer Education and Cognitive Systems (CECS) 1100 classes (Introduction to Computers in Education), students enrolled in graduate CECS classes, and students in the doctoral program in Information Science at the University of North Texas. Participants were selected from various classes during the fall semester of 1997. Participation was voluntary. A total of 110 students participated in the study.

Instrumentation

Demographic Survey

Demographic data were collected for all participants before they began the research study. In addition to the standard demographic data, the questionnaire assessed the level of experience with computers, with the World Wide Web, with Web-based search systems, and with 3-D game programs. Not all demographic data were used in answering the research hypotheses. The Demographic Survey was used to collect the following data: age, gender, education level, experience with computers and the World Wide Web, experience with Web-based search systems, and experience with 3-D game programs.

The Demographic Survey is included in Appendix A.

User Satisfaction Questionnaire

The User Satisfaction Questionnaire was based on the IBM computer usability satisfaction questionnaire developed by James Lewis [, 1995 #15]. This questionnaire is based on a number of usability questionnaires used by human factors engineers, both inside and outside IBM, that have been tested for reliability and validity. The reliability and validity analyses in Lewis's [, 1995 #15] study indicated the following:

> Coefficient alpha analyses showed that the reliability of the overall summative scale (OVERALL) was .97. Therefore, the overall scale [has] excellent reliability. Correlation analyses support the validity of the scales. Using data from the 22 participants who completed all ASQ and all PSSUQ items, the OVERALL scale correlated highly with the sum of the ASQ ratings that participants gave after completing each scenario, r(20) = .80, p = .0001. Using data from the 31 participants who completed all scenarios and all items on the PSSUQ, OVERALL also correlated significantly with the percentage of successful scenario completion, r(29) = -.40, p = .026. (p. 66)

The User Satisfaction Questionnaire contained 12 questions that measured participants' satisfaction with different facets of each search interface. An overall user satisfaction score was obtained by summing the answers to all 12 questions for each search interface. The User Satisfaction Questionnaire was administered to each participant at the end of the test. It contained two parts, one for the 3-D search interface and one for the 1-D search interface. For each question, participants indicated their satisfaction with each search interface by marking on a scale from 1 to 100, with 1 indicating very dissatisfied and 100 indicating very satisfied. The User Satisfaction Questionnaire contained questions that enabled the collection of the following data for each interface (a) how satisfied the participant was with the ease of use of the search interface, (b) how easy the participant thought the search interface was to use, (c) how effectively the participant thought the search interface allowed him/her to complete all the tasks, (d) how quickly the participant thought the search interface allowed him/her to complete all the tasks, (e) how comfortable the participant felt using the search interface, (f) how productive the participant felt using the search interface, (g) how well the participant liked to use the search interface, (h) how satisfied the participant was with the capabilities of the search interface, and (i) the overall satisfaction of the participant with the search interface.

The User Satisfaction Questionnaire is included in Appendix B.

User Preference Questionnaire

Once the participants completed the User Satisfaction Questionnaire, a final questionnaire was administered to determine which interface the participants preferred. Each participant was asked to indicate a preference for either the 3-D search interface or the 1-D search interface. Because all participants used both interfaces to complete the tasks, the participants were able to compare the two interfaces and choose the one they preferred. Participants were then asked to gauge by how much they preferred the chosen interface over the other. Scores could range from 1 to 100, with 1 indicating a very slight preference for the interface and 100 indicating a very strong preference for the interface.

The User Preference Questionnaire is included in Appendix C.

Design and Procedures

The study was designed to test the usability of a three-dimensional (3-D), VRML search interface versus a one-dimensional (1-D), menu-based search interface on the World Wide Web. Participants used both the 3-D search interface and the 1-D search interface to execute predefined search tasks.

Participants were given a Request for Participation form (Appendix E) that described the experiment and were asked to sign an Informed Consent form (Appendix E) for participation in the experiment. A demographic questionnaire was administered to each participant before the actual test began.

Participants were brought into the test situation and asked to execute task scenarios using the two different search interfaces. Data were collected to determine (a) how long it took to complete each task (with a time limit of 60 minutes for the entire test), (b) the participant's satisfaction with each interface, (c) the participant's preference for a search interface, and (d) the participant's comparison of the two search interfaces.

The tasks consisted of four different searches, using either the 3-D search interface or the 1-D search interface. Every participant was asked to complete all four searches, two with the 3-D search interface and two with the 1-D search interface. Search time was recorded for each search that the participant conducted. Participants were allowed to take as much time as they needed for each individual search; however, the total test time was limited to 60 minutes. Test data were excluded from analysis for those participants who could not complete all four searches within the 60-minute time limit.

In order to minimize the effects of search order, participants were randomly assigned to different starting sequences of the task activities based on the following criteria: 3-D or 1-D search interface, Category search or Document search, and Information A or Information B

Two Category searches and two Document searches were conducted, using both the 3-D and 1-D search interfaces by all participants. The Category search using the 3-D and 1-D search interfaces required looking for Information A with one search interface and looking for Information B with the other search interface. Likewise, the Document search using the 3-D and 1-D search interfaces required looking for Information A and

Information B with each search interface. A participant therefore conducted two Web-based searches using the 3-D search interface and two Web-based searches using the 1-D search interface. Appendix D contains the specific items that comprised Information A and Information B for the Category and Document searches.

Category and Document searches were paired with 3-D and 1-D search interfaces, with the pairing randomized throughout the test. Assignment of the search for Information A or Information B was randomized for each Category or Document search. Consequently, 16 different sequences of task activities were generated from these three criteria, to which participants were randomly assigned. The study was designed so that 3-D and 1-D searches were interleaved, with a 3-D search followed by a 1-D search and vice versa. For example, a participant might be presented with searches in this order:

1. Use the 1-D search interface to perform a Category search for Information A.
2. Use the 3-D search interface to perform a Category search for Information B.
3. Use the 1-D search interface to perform a Document search for Information B.
4. Use the 1-D search interface to perform a Document search for Information A.

The order of presentation of these four different tasks were randomized so that search order was not a factor in the study.

Once participants completed all four searches, or when time had expired, they completed the User Satisfaction Questionnaire. Participants then completed the User Preference Questionnaire. The data for a given participant were complete if they finished all tasks and answered all the questions on both the User Satisfaction Questionnaire and the User Preference Questionnaire. A total of 75 participants had complete data out of the 110 who originally participated in the study.

Null Hypotheses

The following null hypotheses were stated for investigation in this study:

1. There is no significant difference in the amount of time required by the same participant to find the specific information using a 3-D search interface and a 1-D search interface.

2. There is no significant difference in user satisfaction scores of the same participant after using a 3-D search interface and a 1-D search interface.

3. There is no significant difference in the percentage of participants preferring a 3-D search interface versus a 1-D search interface.

4. There is no significant difference in user preference scores for participants preferring a 3-D search interface versus a 1-D search interface.

Data Analysis

All statistical analyses of the data were performed using Data Desk®, version 5.0.1 for Macintosh, except for the Cronbach alpha test for internal consistency reliability of the User Satisfaction Questionnaire, which was performed using SPSS®, version 7.0 for Windows. All statistical data were analyzed for two-tailed significance at a 0.01 alpha level. A Cronbach alpha internal consistency reliability was computed for the User Satisfaction Questionnaire scores, with a separate analysis performed for 3-D search interface data and for 1-D search interface data.

Hypotheses 1 and 2 were tested using a dependent or paired t-test. Every participant completed two searches using the 3-D search interface and two searches using the 1-D search interface. Every participant then completed a User Satisfaction Questionnaire for both the 3-D and 1-D search interface. Because every participant was tested on both search interfaces, the data were paired for the analysis. The independent variable for both hypotheses 1 and 2 was search interface. The dependent variable for hypothesis 1 was search time, measured in minutes. The dependent variable for hypothesis 2 was user satisfaction score, measured on a scale from 1 to 100 (1 = very dissatisfied with the search interface to 100 = very satisfied with the search interface).

Hypothesis 3 was tested using a z-test for difference in proportions for independent samples. Participants were asked to indicate a preference for either the 3-D search interface or the 1-D search interface. The proportions of participants who chose either the 3-D search interface or the 1-D search interface were then compared in the data analysis.

Hypothesis 4 was tested using a pooled t-test of the difference between independent means. Participants were asked to indicate by how much they preferred their chosen search interface. The independent variable for hypothesis 4 was search interface, and the dependent variable was user preference score, measured on a scale from 1 to 100 (1 = very slight preference for the chosen search interface to 100 = very strong preference for the chosen search interface). Two participants indicated no preference, and their data were excluded from further analysis.

CHAPTER 4

ANALYSIS OF DATA

Demographic Characteristics of ParticipantsThe age categories of undergraduate and graduate students who participated in the study are presented in Table 1. More than half of the students were between the ages of 20 and 34, which is typical of university students.

Table 1

Age Categories of Participants in the Sample

Age grouping	Frequency	Percent
Less than 20 years old	8	10.7
20 - 24 years old	17	22.6
25 - 29 years old	14	18.7
30 - 34 years old	11	14.7
35 - 39 years old	4	5.3
40 - 44 years old	12	16.0
45 - 49 years old	5	6.7
50 - 54 years old	4	5.3
55 or older	0	0.0
Total	75	100.0

Table 2 presents the breakdown of males and females in the study. Slightly more females than males participated in the study. Table 3 presents the computer experience of the students who participated in the study. Participants typically had 2 or more years of computer experience. Table 4 indicates the amount of World Wide Web experience of the participants. More than half the participants had 1 or more years experience, and almost 75% of the participants had at least 6 months or more experience. Table 5 indicates the amount of experience participants had with some kind of World Wide Web search engine. Almost half the participants had 1 or more years experience, and more than 70% of the participants had at least 6 months or more experience.

Table 2

Gender of Participants in the Sample

Gender	Frequency	Percent
Male	32	42.7
Female	42	56.0
Missing	1	1.3
Total	75	100.0

Table 3

Computer Experience of Participants in the Sample

Computer experience	Frequency	Percent
1 week or less	1	1.3
1 week to 1 month	0	0.0
1 month to 6 months	2	2.7
6 months to 1 year	5	6.7
1 year to 2 years	12	16.0
2 years or more	55	73.3
Total	75	100.0

Table 4

World Wide Web Experience of Participants in the Sample

World Wide Web experience	Frequency	Percent
1 week or less	3	4.0
1 week to 1 month	5	6.7
1 month to 6 months	11	14.6
6 months to 1 year	18	24.0
1 year to 2 years	15	20.0
2 years or more	23	30.7
Total	75	100.0

Table 5

Search Engine Experience of Participants in the Sample

Search engine experience	Frequency	Percent
No experience	2	2.7
1 week or less	4	5.3
1 week to 1 month	4	5.3
1 month to 6 months	11	14.7
6 months to 1 year	19	25.3
1 year to 2 years	17	22.7
2 years or more	18	24.0
Total	75	100.0

Reliability of the User Satisfaction Questionnaire

A Cronbach alpha internal consistency reliability coefficient was computed for participants' scores on the User Satisfaction Questionnaire for the 3-D search interface and for the 1-D search interface. There were 75 participants and 12 items in the questionnaire. The Cronbach alpha coefficient for question responses for the 3-D search interface was 0.96. The Cronbach alpha coefficient for question responses for the 1-D search interface was 0.95. These results indicate a very high internal consistency of participants' responses to the questions dealing with both the 3-D search interface and the 1-D search interface.

Hypothesis 1

The first hypothesis tests for no significant difference in the time required to find information using either a 3-D search interface or a 1-D search interface by the same participants. The amount of time for each participant to find the information in the 1-D and 3-D interfaces was summed and is presented in minutes. Because a participant conducted searches using both the 3-D search interface and the 1-D search interface, a dependent or paired t-test was calculated. The average search time for each interface, standard deviations, and t-test calculations are presented in Table 6.

Table 6

Interface Means, Standard Deviations, and Paired t-test for Search Time

Interface	Mean (minutes)	SD	Mean difference	Standard error	t	df	p
3-D	14.97	8.55					
			9.07	1.12	8.09	74	$p \leq 0.0001$
1-D	5.90	4.34					

The null hypothesis of no difference in search time is rejected, and the alternative hypothesis is accepted. There was a significant difference in the amount of time it took participants to find information using a 3-D and 1-D search interface. Participants took less time to find information using the 1-D search interface versus the 3-D search interface.

Hypothesis 2

The second hypothesis tests for no significant difference in user satisfaction scores after conducting searches using both a 3-D search interface and a 1-D search interface. Participants completed a User Satisfaction Questionnaire after conducting a search for information using both the 3-D and the 1-D search interfaces. The User Satisfaction Questionnaire contained 12 questions, and scores could range from 1 to 100 (1 = strong dissatisfaction with the interface to 100 = strong satisfaction with the interface). See Appendix B for a copy of the questionnaire. A comparison of the 12 user satisfaction means and standard deviations for each question for both 3-D and 1-D search interfaces is presented in Table 7.

Table 7

Means and Standard Deviation for User Satisfaction Questionnaire Questions by Interface

Question		Mean	SD
Question 1	3-D	33.44	29.13
	1-D	77.68	20.28
Question 2	3-D	34.53	28.14
	1-D	78.04	20.15
Question 3	3-D	48.80	32.09
	1-D	79.49	21.32

Question 4	3-D	33.39	29.19
	1-D	75.25	23.54
Question 5	3-D	38.36	30.99
	1-D	74.67	22.80
Question 6	3-D	35.40	32.24
	1-D	82.33	20.76
Question 7	3-D	54.12	30.80
	1-D	84.84	17.76
Question 8	3-D	6.88	32.56
	1-D	83.89	17.89
Question 9	3-D	42.73	33.48
	1-D	74.15	23.20
Question 10	3-D	43.21	33.22
	1-D	75.76	20.64
Question 11	3-D	45.21	30.06
	1-D	74.25	20.74
Question 12	3-D	41.35	32.83
	1-D	77.21	19.61

A participant's user satisfaction score was derived by summing responses to each of the 12 questions. Every participant completed the User Satisfaction Questionnaire for each interface, so two scores were obtained for each participant. Scores could range from 1 to 100 (1 = <u>strong dissatisfaction</u> with the interface to 100 = <u>strong satisfaction</u> with the interface), so 12 was the worst possible score and 1200 was the best possible score. Because each participant conducted searches using both 3-D and 1-D search interfaces, a dependent or paired t-test for differences in participants' satisfaction scores was calculated. Average satisfaction scores for each interface, standard deviations, and t-test calculations are presented in Table 8.

Table 8

Interface Means, Standard Deviations, and Paired t-test for User Satisfaction Scores

Interface	Mean	SD	Mean difference	Standard error	t	df	p
3-D	497.42	313.17					
			440.15	45.47	9.680	74	$p < 0.0001$
1-D	937.57	201.55					

The null hypothesis of no difference in participants' satisfaction scores is rejected, and the alternative hypothesis accepted. There was a significant difference in the user satisfaction scores between the 3-D and 1-D search interfaces. Participants had a higher average level of satisfaction with the 1-D search interface versus the 3-D search interface.

Hypothesis 3

The third hypothesis tests for no significant difference in the percentage of participants who would choose the 3-D search interface versus the 1-D search interface. Participants were asked to state a preference for either the 3-D or the 1-D search interface on the User Preference Questionnaire. Out of 75 participants who had complete data, 2 participants indicated No Preference for either interface. The 2 participants were excluded from the analysis. A z-test for differences in proportions between two groups was calculated. The number of participants in each group, group percentages, and z-test calculations are presented in Table 9.

Table 9

Number of Participants, Group Percentage and z-test for Differences in Proportions between Interface Preference Groups

Interface preference group	N	Percent	z	p
Prefer 3-D	17	23.29		
			4.026	$p < 0.0001$
Prefer 1-D	56	76.71		
Total	73	100.00		

The null hypothesis of no significant difference in the percentage of participants who would choose the 3-D search interface versus the 1-D search interface was rejected, and the alternative hypothesis accepted. There is a significant difference in the percentage of participants who chose the 3-D interface versus the 1-D search interface. More participants selected the 1-D interface.

Hypothesis 4

The fourth hypothesis tests for no significant difference in preference scores between two self-selected groups after searches were conducted using the 3-D and 1-D search interfaces. The preference scores were obtained from a question on the User Preference Questionnaire. After participants had indicated a preference for one interface or the other, they were instructed to indicate by how much they preferred the interface they chose. Scores could range from 1 to 100 (1 = very slight preference for the interface to 100 = very strong preference for the interface). Participants self-selected themselves into one of two possible groups: Prefer 3-D or Prefer 1-D. An independent or pooled t-test was then calculated to test for mean preference score differences between the two groups. Interface preference group means, standard deviations, and pooled t-test calculations are presented in Table 10.

Table 10

Group Means, Standard Deviations, and Pooled t-test for Preference Scores

Group	Mean	SD	Mean difference	Standard error	t	df	p
Prefer 3-D	43.23	20.83					
			22.50	8.02	2.807	71	p = 0.0065
Prefer 1-D	65.73	30.91					

The null hypothesis of no difference in participants' satisfaction scores is rejected, and the alternative hypothesis accepted. There was a significant mean difference in preference scores between those participants who preferred the 3-D search interface and those who preferred the 1-D search interface. Those participants who preferred the 1-D search interface had higher average preference scores than those participants who preferred the 3-D search interface.

CHAPTER 5

SUMMARY OF FINDINGS, CONCLUSIONS, AND RECOMMENDATIONS

Introduction

The purpose of this study was to determine how people would interact with a three-dimensional (3-D) interface to a World Wide Web database versus a conventional text-based, one-dimensional (1-D) interface to the same database. Variables of interest in this study were the time it took to find information in the database, user satisfaction with the search interface as measured by a user satisfaction survey, and user preference for search interface.

The participants for the study consisted of undergraduate and graduate students who currently use World Wide Web search engines, such as Yahoo, to search for information on the Internet, and who use a Web browser with a graphical user interface, such as Netscape Navigator. The sample for this study was drawn from students enrolled in Computer Education and Cognitive Systems (CECS) 1100 classes (Introduction to Computers in Education), students enrolled in graduate CECS classes, and students in the doctoral program in Information Science at the University of North Texas. Participants were selected from various classes during the Fall semester of 1997. Participation was voluntary. A total of 110 students participated in the study. Of this total, 75 participants completed all sections of the study, and their data were retained for analysis.

Test participants were asked to perform four different searches. Two searches used the 3-D search interface, and two searches used the 1-D search interface. Searches consisted of two different Category items and two different Document items. All participants performed the same four searches; however, in order to minimize the effects of search order, the sequence of searches was randomized across the test. Search time was recorded for each search that the participant conducted. Participants were allowed to take as much time as they needed for each individual search; however, the total test time was limited to 60 minutes. Test data were excluded from analysis for those participants who could not complete all four searches within the 60-minute time limit.

Once the participants had performed all four searches, or time had run out for the search portion of the test, participants completed the User Satisfaction Questionnaire. This questionnaire consisted of 12 questions measuring user satisfaction with different aspects of the search interface. For each question, participants indicated their satisfaction with the interface by marking on a scale from 1 to 100, with 1 indicating strongly disagree and 100 indicating strongly agree, with statements regarding their satisfaction with the interface. Each question had two scale lines following it, one for the 1-D interface and one for the 3-D interface, so that participants indicated satisfaction with both interfaces in the same question. An overall score for the user satisfaction was obtained by summing the answers to all 12 questions for each search interface. Summed scores ranged from 12 to 1200.

Following the User Satisfaction Questionnaire, participants completed the User Preference Questionnaire. Participants indicated a preference for either the 3-D search interface or the 1-D search interface. Participants then marked a scale line indicating then amount of preference for the chosen search interface. Scores ranged from 1 to 100, with 1 indicating very slight preference for the interface and 100 indicating very strong preference for the interface.

Summary of Major Findings

The major findings of this study are as follows:

1. The mean search time for the 3-D search interface was 14.97 minutes. The mean search time for the 1-D search interface was 5.90 minutes. A paired t-test

performed on this data confirmed that the difference in the means was significant at less than the 0.0001 level. The null hypothesis that search times for the 3-D search interface would not be significantly different from search times for the 1-D search interface was rejected. The data show that search times for the 1-D search interface were significantly less than search times for the 3-D search interface.

2. The mean user satisfaction score for the 3-D search interface was 497.42. The mean user satisfaction score for the 1-D search interface was 937.57. A paired t-test performed on this data confirmed that the difference in the means was significant at less than the 0.0001 level. The null hypothesis that user satisfaction with the 3-D search interface would not be significantly different from user satisfaction with the 1-D search interface was rejected. The data show that user satisfaction with the 1-D search interface was significantly greater than user satisfaction with the 3-D search interface.

3. Of the 75 participants in the study, 2 indicated no preference for search interface, and their data were excluded from further analysis. The data showed that 17 participants (23.29%) preferred the 3-D search interface and 56 participants (76.71%) preferred the 1-D search interface. A z-test for differences in proportions performed on this data confirmed that the difference in the frequencies was significant at less than the 0.0001 level. The null hypothesis that there would not be a significant difference in the percentage of users who preferred the 3-D search interface versus the 1-D search interface was rejected. The data show that users' preference for the 1-D search interface was significantly greater than users' preference for the 3-D search interface.

4. The mean user preference score for the 3-D search interface was 43.23. The mean user preference score for the 1-D search interface was 65.73. A pooled t-test for independent means performed on this data confirmed that the difference in the means was significant at the 0.0065 level. The null hypothesis that user preference score for the 3-D

search interface would not be significantly different from user preference score for the 1-D search interface was rejected. The data show that the users' preference score for the 1-D search interface was significantly greater than users' preference score for the 3-D search interface.

Discussion of Major Findings

Each part of this research was concerned with discovering whether or not 3-D interfaces would perform as well as or better than 1-D interfaces. In each case, the 3-D interface did not perform as well as the 1-D interface, for time to search as well as user satisfaction with the interfaces.

The time taken to find information using the 3-D interface was more than 2.5 times longer than the time taken to find information using the 1-D interface. Some of this can be attributed to the fact that interface response time in a 3-D VRML environment is slower than interface response time in a 1-D text environment. This is a function of the computer hardware available to the average consumer, and these results will definitely carry over to the real world. However, the average search times for both the 1-D interface and the 3-D interface were much longer than the time it took to load either the 1-D or 3-D interfaces. Users spent much more time searching the interfaces for information than they did waiting for information to appear on the screen. The 3-D environment made it more difficult for many users to navigate through the information to find what they wanted, and this resulted in the much longer search times than with the 1-D interface.

That users found the 1-D interface easier to use and more effective in completing the searching tasks is reflected in the both the user satisfaction scores and in the user interface preference. User satisfaction scores for the 1-D interface were almost twice as high as for the 3-D interface, whereas users preferred the 1-D interface to the 3-D interface by a margin of more than 3 to 1. In addition, the fact that the user preference score for the 1-D interface was 1.5 times greater than the user preference score for the 3-D interface indicates that even those users who preferred the 3-D interface were not as positive in their support for the 3-D interface as were those users who preferred the 1-D interface. In all measures used in this study, the 1-D interface was the interface of choice for simple searching tasks.

Conclusions

Based on the methodology and findings of this study, the following is concluded:

1. Users need to be careful when evaluating 3-D interfaces based on the current publicity given to them. Although most popular discussions of 3-D interfaces are extremely positive, this research shows that simple search tasks using a hierarchical 3-D search interface can be slower and more difficult to perform than the same search tasks using a 1-D search interface.

2. Current consumer-level computer technology is not adequate to support 3-D VRML interfaces of any complexity. Even though the search interface used in this study was relatively simple and the computers used for the test were high-end consumer-level machines, the performance of the 3-D interface was not acceptable to most people.

3. Using 3-D interfaces and objects to represent nonphysical concepts and ideas may not present an advantage for users versus 1-D representations of the same concepts and ideas.

Discussion of Conclusions

In both the popular press [Larijani, 1994 #68; Broll, 1996 #6; Chinnock, 1996 #21; Flohr, 1996 #13; Pimental, 1993 #2; Rheingold, 1991 #1; Pesce, 1995 #10] and in scholarly writing [Ames, 1996 #4; Darken, 1996 #29; Deering, 1995 #34; Ellis, 1994 #23; Gallimore, 1993 #35; Latta, 1994 #80; Ribarsky, 1994 #82; Stuart, 1996 #69;

Waterworth, 1994 #24], there are few cautionary words about the use of VR in any kind of searching environment. When problems are mentioned with VR environments, they are usually about simulator sickness [Strauss, 1995 #27] or about the problems people have adjusting to the real world after being in a VR environment [Krueger, 1996 #26]. The implications of currently published articles is that if users try a 3-D interface and use it to perform a real world task, that they will prefer it to the standard interfaces they more commonly use. The results of this study show that there is a need for further study to determine just where VR interfaces are appropriate, where they will be most effective, and whether or not people will want to use them when and if they are deployed.

It takes more time to plan, design, and program a 3-D interface than a simple 1-D text-based menu system. In addition, programming a 3-D interface requires special programming skills that few people possess. VRML makes the process of programming in 3-D easier than using a general purpose programming language such as C or C++, but it still requires sophisticated programming skills. VRML interface builders make the process easier and faster, but it is still more difficult and time-consuming to build a 3-D interface compared to a 1-D interface. Finally, 3-D interfaces require faster, more expensive communications devices and faster, more powerful, and more expensive machines for the user to be able to display and run the 3-D interfaces. The results of this test show that it may not be worth the extra time and money required to build a 3-D search interface, for simple searches of conventional databases, when the 1-D interface may perform better, be better liked, and be easier and less expensive to build.

The usability of software is generally dependent on a number of factors, including (a) speed of task completion while using the interface, (b) the number of errors made while completing a task, (c) how comfortable a user feels while completing a task, (d) how simple the user feels the interface is to master, (e) how pleasant the interface feels to the user, and (f) how much function the user feels the interface contains [Lewis, 1995 #15; Shneiderman, 1980 #93]. The time taken to find an item using the search interfaces and the user satisfaction questionnaire that was administered to participants in this study attempted to gauge the participants' satisfaction with the search interfaces on all these measures. Both the time to search measure and the user satisfaction measure indicate that the 3-D interface in this study was not as usable as the 1-D interface.

People use search engines on the World Wide Web to help them find information more quickly than simply wandering from Web site to Web site. Acceptance or rejection of a search interface is directly related to how fast it lets the user find information [Marchionini, 1987 #87; Marchionini, 1995 #91; Janes, 1996 #8]. The fact that the 1-D interface allowed users to find information more than 2.5 times as fast as the 3-D interface means that the 1-D interface was a better tool for searching the World Wide Web. It is not surprising then, that when users were asked directly which interface they preferred, they chose the 1-D interface by a margin of almost 2 to 1.

Most research studies use high-powered computers as their test machines. The computer systems are either special-purpose, 3-D computers [Rheingold, 1991 #1], or are fast general-purpose, 3-D computers [Ware, 1993 #67]. In either case, the machines used for most VR testing are not the types of systems that most consumers can afford or will be likely to afford for a number of years. Any research that concludes that VR systems will make it easier for the average computer user to do things should take into account the level of computer hardware available to the average person.

Examples of uses for VR given in the popular press and in the literature include virtual travel, 3-D games, telepresence, telesurgery, positioning radiation beams for treatment, architectural design, games, virtual theater, virtual art museum, virtual classroom, 3-D participant lessons, dangerous lab experiments, armchair tourism, flight simulation, battle planning, erotica, study of insanity, and fantasy [Larijani, 1994 #68]; outdoor landscape design, artistic designs, business money flow, stock market data analysis, chemical and molecular simulation, virtual prototyping, virtual bodies, and auto test driving [Pimental, 1993 #2]; and entertainment, virtual humans, molecular modeling, concept design, kitchen design, and virtual manufacturing [Harrison, 1996 #70]. Most of the systems in this wish list represent models of real-world places and objects. Current research in VR systems also uses real-world places or objects as models for the 3-D systems [Darken, 1996 #29; Dede, 1994 #73; Deering, 1995 #34; Gallimore, 1993 #35; Ware, 1993 #67; Ware, 1996 #83]. This research focused on representing concepts and ideas that do not have a direct physical representation. For such systems, there is no clear advantage in the use of 3-D interfaces to model the systems.

Recommendations and Implications

One of the design points of this study was to use consumer-level machines in the design and testing of the 3-D interface. Computer technology is advancing quickly, and what seems fast today will be slow 20 years from now, and what is high-end today will be consumer-level tomorrow. In order to test 3-D versus 1-D for the machines of the future, the research should be repeated with very fast, expensive machines to see if the speed of response of the interface will change the result.

When doing research of this type, the researcher should make sure all participants have adequate search skills. The intent of a study of this kind is to test the adequacy of the search interface and not the adequacy of the participants' search skills. In this research, the data from several participants had to be discarded because the participants could not complete a single search.

The researcher should provide more training time on both the 3-D search interface and the 1-D search interface. Again, the intent of this research was to test the efficiency and user satisfaction with the search interface, not how fast people could learn to use an unfamiliar computer interface. The use of the 1-D and 3-D search interfaces was intermixed in the study, but some of the differences in search time and user satisfaction could have been due to users' unfamiliarity with the 3-D systems in general and with 3-D search interfaces versus 1-D interfaces.

The researcher should capture and do an analysis on the specific pathways that users follow to find information using the 1-D and 3-D search interfaces. Is there a difference in the search pathways that users follow when using 1-D and 3-D interfaces? Do users take more or fewer blind alleys in their data search with 1-D or 3-D? Are the times for 1-D faster because the search strategy is more effective with 1-D versus 3-D?

In this study, there was a significant difference between Category and Document searches. Expectations at the beginning of the study were that there would be no difference between Category and Document searches. Is this a real difference or just an artifact of the test? A further study should be conducted that matches Category and Document searches and tests to determine whether or not there is still a significant difference in the search times.

APPENDIX A

DEMOGRAPHIC SURVEY

ID #__________ **Demographic Survey** Comp #___________

All responses to this survey are kept confidential, please Do Not write your name on the form.

This section of the survey gathers general information about you and your knowledge of computers, the Internet and the World Wide Web.

Please check the blank which applies to you or fill in the information requested.

1. Age:
   ( ) less than 20 ( ) 20 - 24 ( ) 25 - 29
   ( ) 30 - 34 ( ) 35 - 39 ( ) 40 - 45
   ( ) 45 - 49 ( ) 50 - 54 ( ) 55 or older

2. Gender: ( ) Male ( ) Female

3. Education Level Completed: ( ) High School ( ) Some College ( ) Bachelors degree ( ) Masters degree ( ) Doctorate or professional degree

4. How much experience with computers do you have?

   ( ) 1 week or less ( ) 1 week to 1 month ( ) 1 month to 6 months
   ( ) 6 months to 1 year ( ) 1 year to 2 years ( ) 2 years or more

   4a. How many hours per week do you use a computer? ____________

5. How much experience with the World Wide Web do you have?

   ( ) 1 week or less ( ) 1 week to 1 month ( ) 1 month to 6 months
   ( ) 6 months to 1 year ( ) 1 year to 2 years ( ) 2 years or more

   5a. How many hours per week do you use the World Wide Web? ____________

6. Have you ever used a World Wide Web search engine such as Yahoo?

   ( ) Yes ( ) No

   6a. If yes, how much experience with World Wide Web search engines do you have:

   ( ) 1 week or less ( ) 1 week to 1 month ( ) 1 month to 6 months
   ( ) 6 months to 1 year ( ) 1 year to 2 years ( ) 2 years or more

   6b. How many hours per week do you use World Wide Web search engines? ____________

   6c. Yahoo has both a search box and a system of menus. When you use a search engine, what percentage of the time do you use the menu system? ____________

7. Do you play 3-D computer games?

( ) Yes ( ) No

7a. If yes, how many hours per week do you play 3-D computer games? ____________

APPENDIX B

USER SATISFACTION QUESTIONNAIRE

ID#__________ **User Satisfaction Questionnaire**

This questionnaire gives you an opportunity to tell us your reactions to the two search application you used. Your responses will help us understand what aspects of the applications you are particularly concerned about and the aspects that satisfy you. To as great a degree as possible, think about all the tasks that you have done with the two applications while you answer these questions. Please read each statement and indicate how strongly you agree or disagree with the statement by making a mark on the scale next to the search application.

When answering the questions, please keep in mind your experience with both the 3-D and the 1-D search interface.

NOTE: Each question asks about your experience with both the 3-D and the 1-D search interface.

Thank you!

1. Overall, I am satisfied with how easy it is to use this search application.

strongly disagree ... strongly agree

3-D Interface | -- |

1-D Interface | -- |

2. It was simple to use this search application.

strongly disagree ... strongly agree

3-D Interface | -- |

1-D Interface	| -- |

3. I could effectively complete the tasks using this search application.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

4. I was able to complete the tasks quickly using this search application.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

5. I was able to efficiently complete the tasks using this search application.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

6. I felt comfortable using this search application.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

7. It was easy to learn to use this search application.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

8. I believe I could become productive quickly using this search application.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

9. The interface of this search application was pleasant.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

10. I liked using the interface of this search application.

strongly strongly

	disagree	agree
3-D Interface	| --	|
1-D Interface	| --	|

11. This search application has all the functions and capabilities I expect it to have.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

12. Overall, I am satisfied with this search application.

	strongly disagree		strongly agree
3-D Interface	|	--	|
1-D Interface	|	--	|

APPENDIX C

USER PREFERENCE QUESTIONNAIRE

ID#__________ **User Preference Questionnaire**

This final questionnaire gives you the opportunity to tell us which interface you prefer. Please place a **single** mark on the line below to indicate the amount of your preference for one interface over the other, where 100 indicates a complete preference for one interface over the other and 0 indicates no preference for one or the other interface.

A mark to the left of the 0 indicates a preference for the 1-D interface. A mark to the right of the 0 indicates a preference for the 3-D interface.

Prefer **1-D**	No Preference	Prefer **3-D**
|---	|	---|
100	0	100

Thank you!

APPENDIX D

THE TEST

Search Instructions

Each of the following pages contains instructions for performing a single search.

When the instructions ask you to find a document, you are being asked to find a single Web page which contains the specific information.

When the instructions ask you to find a category, you are being asked to find a Yahoo category, which may contain a number of documents and sub-categories.

Click on **Button A**.

Turn to the next page.

Document Search

Find the document, labeled Simpson, Alan D., which contains the resume of Alan D. Simpson.

Click on **Button B**.

Turn to the next page.

Document Search

Find the document, labeled California Student Aid Commission, which contains information about California college grants, scholarships, and student loans for the whole California area.

Click on **Button A**.

Turn to the next page.

Category Search

Find the category which contains information about careers in Broadcasting.

Click on **Button B**.

Turn to the next page.

Category Search

Find the category which contains information about teaching English as a second language.

Click on **Button A**.

Turn to the next page.

End of Test

Congratulations! You are done.

Please complete the forms on the following pages.

APPENDIX E

REQUEST FOR PARTICIPATION

AND

USE OF HUMAN SUBJECTS CONSENT FORM

REQUEST FOR YOUR PARTICIPATION IN A RESEARCH STUDY

3-D Interface to the World Wide Web

Dear CECS 1100 Student

This is a request for your participation in a research study of CECS 1100 students. The study is being conducted for doctoral dissertation research by Peter Scannell. You are being asked to participate in an experiment to test the effectiveness of a three-dimensional interface to Yahoo, on the World Wide Web. The experiment is designed to take about half-an-hour and will involve your searching for documents and categories on both the conventional Yahoo interface and a 3-D interface to Yahoo. This will give you a chance to experience 3-D interaction on the World Wide Web.

It is possible that some people may experience some disorientation while navigating through a 3-D world on the computer. If this happens to you, you can stop the experiment at any time. Your participation is entirely voluntary and I will not be keeping a record of your name or class section. Again, if at any time during the experiment you become uncomfortable with the tasks you are performing or navigating through the 3-D world, and want to quit, you can do so at any time without penalty.

Participation in this study will require you to:

- complete a demographic survey form, which will be identified with an ID # only.
- search for documents and categories using both the standard Yahoo interface and the 3-D Yahoo interface.
- complete a questionnaire after the interfaces.
- complete a user preference questionnaire.

In each case, the only identifying marks on the questionnaires or experiment results will be the ID #, which will not be linked to your name or class section. Al forms will be kept confidential and you will not be identified as an individual in any data analysis.

Participation in this study in no way affects your grade in this course. You are not required to participate in this study. You may withdraw you consent at any time without penalty.

USE OF HUMAN SUBJECTS
INFORMED CONSENT

NAME OF SUBJECT:____

I have heard a clear explanation of the research study in which I am being asked to participate. I understand that this research is being used as a dissertation for Peter Scannell, doctoral student in Information Science.

I agree to:

- complete a demographic survey form,
- search for documents and categories using both the standard Yahoo interface and the 3-D Yahoo interface.
- complete a questionnaire after using the interfaces.
- complete a user preference questionnaire.

In addition, I give my consent for my scores on the search, questionnaire, and user preference questionnaire to be compared with other scores of participants in the study. I understand that these scores will be kept confidential an I will not be identified as an individual. **I understand that I may withdraw my consent at any time without penalty.**

With my understanding of this, having received information about the study and satisfactory answers to the questions I may have asked, I voluntarily consent to participate in this study. If I have any questions regarding this research, I may contact Peter Scannell at (972) 492-6517.

SIGNED:____ ____________________

SUBJECT DATE

SIGNED:____ ____________________

WITNESS DATE

This project has been reviewed and approved by the University of North Texas Committee for Protection of Human Subjects

REFERENCES

Ahlberg, C., & Shneiderman, B. (1994a). The Alphaslider: A compact and rapid selector. In B. Adelson, S. Dumais, & J. Olson (Eds.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systemspp. 365-371). New York: ACM.

Ahlberg, C., & Shneiderman, B. (1994b). Visual information seeking: Tight coupling of dynamic query filters with Starfield displays. In B. Adelson, S. Dumais, & J. Olson (Eds.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systemspp. 313-317). New York: ACM.

Ahlberg, C., Williamson, C., & Shneiderman, B. (1992). Dynamic Queries for information exploration: An implementation and evaluation. In P. Bauersfeld, J. Bennett, & G. Lynch (Eds.), Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systemspp. 619-626). New York: ACM.

Amdur, D. (1995, November). The scene is set for multimedia on the Web. New Media, 42-52.

Ames, S. L., Nadeau, D. R., & Moreland, J. L. (1996). The VRML sourcebook. New York: John Wiley & Sons, Inc.

Berners-Lee, T. (1996,). Longer bio for Tim Berners-Lee [WWW document]. URL http://www.w3.org/pub/WWW/People/Berners-Lee-Bio.html/Longer.html

Bolter, J., Hodges, L. F., Meyer, T., & Nichols, A. (1995). Integrating Perceptual and Symbolic Information in VR. IEEE Computer Graphics and Applications, 15(4), 8-11.

Broll, W., & Koop, T. (1996). VRML: Today and tomorrow. Computers & Graphics, 20(3), 427-434.

Brown, M. H., Meehan, J. R., & Sarkar, M. (1993). Browsing graphs using a fisheye view. In S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel, & T. White (Eds.), Proceedings of ACM INTERCHI'93 and CHI'93 Conference on Human Factors in Computing Systemspp. 516). New York: ACM.

Card, S. K., Pirolli, P., & Mackinlay, J. D. (1991). The cost-of-knowledge characteristic function: Display evaluation for direct-walk dynamic information visualizations. In S. P. Robertson, G. M. Olson, & J. S. Olson (Eds.), Proceedings of ACM CHI’91 Conference on Human Factors in Computing Systemspp. 238-244). New York: ACM.

Card, S. K., Robertson, G. G., & Mackinlay, J. D. (1991). The Information Visualizer, an information workspace. In S. P. Robertson, G. M. Olson, & J. S. Olson (Eds.), Proceedings of ACM CHI’91 Conference on Human Factors in Computing Systemspp. 181-188). New York: ACM.

Chalmers, M., & Chitson, P. (1992). Bead: Explorations in information visualization, Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrievalpp. 330-337). London: Springer-Verlag.

Chang, S.-J., & Rice, R. E. (1993). Browsing: a multidimensional framework. In M. E. Williams (Ed.), Annual Review of Information Science and Technology (Vol. 28, pp. 231-277). Medford, NJ: Learned Information, Inc.

Chimera, R. (1992). Value Bars: An information visualization and navigation tool for multi-attribute listings. In P. Bauersfeld, J. Bennett, & G. Lynch (Eds.), Proceedings of ACM CHI’92 Conference on Human Factors in Computing Systemspp. 293-294). New York: ACM.

Chinnock, C. (1996, March). Virtual reality goes to work. Byte, 26-27.

Darken, R. P., & Sibert, J. L. (1996). Navigating large virtual spaces. International Journal of Human-Computer Interaction, 8(1), 49-71.

Dede, C. J., Salzman, M., & Loftin, R. B. (1994). The development of a virtual world for learning Newtonian mechanics. In P. Brusilovsky, P. Kommers, & N. Streitz (Eds.), Multimedia, hypermedia, and virtual reality: models, systems, and applications; first international conference, MHVR'94, Moscow, Russia, September 14-16, 1994, selected papers (Vol. 1077, pp. 87-106). Berlin: Springer.

Deering, M. F. (1995). HoloSketch: A virtual reality sketching/animation tool. ACM Transactions on Computer-Human Interaction, 2(3), 220-238.

Ellis, S. R. (1994). What are virtual environments? IEEE Computer Graphics & Applications, 14(1), 17-22.

Encarnacao, J., Foley, J., Bryson, S., Feiner, S. K., & Gershon, N. (1994). Research issues in perception and user interfaces. IEEE Computer Graphics and Applications, 14(2), 67-69.

Flohr, U. (1996, March). Put the space in cyberspace. Byte, 61-64.

Fowler, R. H., Fowler, W. A. L., & Wilson, B. A. (1991). Integrating query, thesaurus, and documents through a common visual representation, Proceedings of the 14th Annual International ACM SIGIR Conference on Research and Development in Information Retrievalpp. 142-151). London: Springer-Verlag.

Furnas, G. W. (1986). Generalized fisheye views. In M. Mantei & P. Orbeton (Eds.), Proceedings of ACM CHI'86 Conference on Human Factors in Computing Systemspp. 16-23). New York: ACM.

Gallimore, J. J., & Brown, M. E. (1993). Visualization of 3-D computer-aided design objects. International Journal of Human-Computer Interaction, 5(4), 361-382.

Globus, A., & Uselton, S. (1995). Evaluation of visualization software. Computer Graphics, 29(2), 41-44.

Hand, C. (1994). Other faces of virtual reality. In P. Brusilovsky, P. Kommers, & N. Streitz (Eds.), Multimedia, hypermedia, and virtual reality: models, systems, and applications; first international conference, MHVR'94, Moscow, Russia, September 14-16, 1994, selected papers (Vol. 1077, pp. 107-116). Berlin: Springer.

Harrison, D., & Jaques, M. (1996). Experiments in virtual reality. Oxford: Butterworth-Heinemann.

Hemmje, M., Kunkel, C., & Willett, A. (1994,). LyberWorld - A visualization user interface supporting fulltext retrieval. Paper presented at the Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, London.

Hughes, K. (1994, May). Entering the World-Wide Web: A guide to cyberspace [WWW document]. URL http://www.eit.com/web/.www.guide/guide.04.html

Jacob, R. J. K., Sibert, L. E., McFarlane, D. C., & Mullen, M. P., Jr. (1994). Integrality and separability of input devices. ACM Transactions on Computer-Human Interaction, 1(1), 3-26.

Janes, J. W., & Rosenfeld, L. B. (1996). Networked information retrieval and organization: Issues and questions. Journal of the American Society for Information Science, 47(9), 711-715.

Johnson, B. (1992). TreeViz: Treemap visualization of hierarchically structured information. In P. Bauersfeld, J. Bennett, & G. Lynch (Eds.), Proceedings of ACM CHI’92 Conference on Human Factors in Computing Systemspp. 369-370). New York: ACM.

Johnson, B., & Shneiderman, B. (1991). Tree-Maps: A space-filling approach to visualization of hierarchical information structures, Proceedings of IEEE Visualization '91pp. 284-291). San Diego, CA: IEEE Computer Society Press.

Keim, D. A., & Kriegel, H.-P. (1994). VisDB: Database exploration using multidimensional visualization. IEEE Computer Graphics and Applications, 14(5), 40-49.

Keim, D. A., Kriegel, H.-P., & Seidl, T. (1993). Visual feedback in querying large databases, Proceedings of Visualization '93pp. 158-165). Los Alamitos, CA: IEEE Computer Society Press.

Kjelldahl, L., & Prime, M. (1995). A study on how depth perception is affected by different presentation methods of 3D objects on 2D display. Computers and Graphics, 19(2), 199-202.

Korfhage, R. R. (1991). To see, or not to see - Is that the query?, Proceedings of the 14th Annual International ACM SIGIR Conference on Research and Development in Information Retrievalpp. 134-141). London: Springer-Verlag.

Krol, E., & Klopfenstein, B. C. (1996). The whole Internet user's guide & catalog. Belmont, CA: Wadsworth Publishing Company.

Krueger, M. W. (1996). Forward: The interface will not disappear. International Journal of Human-Computer Interaction, 6(1), 1-3.

Kumar, H. P., Plaisant, C., & Shneiderman, B. (1997). Browsing hierarchical data with multi-level dynamic queries and pruning. International Journal of Human-Computer Studies, 46, 103-124.

Larijani, L. C. (1994). The virtual reality primer. New York: McGraw-Hill, Inc.

Latta, J. N., & Oberg, D. J. (1994). A conceptual virtual reality model. IEEE Computer Graphics & Applications, 14(1), 23-29.

Lewis, J. R. (1995). IBM computer usability satisfaction questionnaires: Psychometric evaluation and instructions for user. International Journal of Human-computer Interaction, 7(1), 57-78.

Lin, X. (1997). Map displays for information retrieval. Journal of the American Society for Information Science, 48(1), 40-54.

Machover, C., & Tice, S. E. (1994). Virtual Reality. IEEE Computer Graphics & Applications, 14(1), 15-16.

Mackinlay, J. D., Robertson, G. G., & Card, S. K. (1991). The Perspective Wall: Detail and context smoothly integrated. In S. P. Robertson, G. M. Olson, & J. S. Olson (Eds.), Proceedings of ACM CHI'91 Conference on Human Factors in Computing Systemspp. 173-179). New York: ACM.

Marchionini, G. (1987). An invitation to browse: designing full-text systems for novice users. The Canadian Journal of Information Science, 12(3), 69-79.

Marchionini, G. (1989). Information-seeking strategies of novices using a full-text electronic encyclopedia. Journal of the American Society for Information Science, 40(1), 54-66.

Marchionini, G. (1995). Information seeking in electronic environments. New York: Cambridge University Press.

Olsen, K. A., Korfhage, R. R., Sochats, K. M., Spring, M. B., & Williams, J. G. (1993). Visualization of a document collection: The VIBE system. Information Processing & Management, 29(1), 69-81.

Pedersen, G. S. (1993). A browser for bibliographic information retrieval, based on an application of lattice theory, Proceedings of the 16th Annual International ACM SIGIR Conference on Research and Development in Information Retrievalpp. 270-279). London: Springer-Verlag.

Pesce, M. (1995). VRML - Browsing and building cyberspace. Indianapolis, IN: New Riders Publishing.

Pimental, K., & Teixeira, K. (1993). Virtual reality: Through the new looking glass. New York: Intel/Windcrest/McGraw-Hill, Inc.

Plaisant, C., & Jain, V. (1994). Dynamaps: Dynamic queries on a health statistics atlas. In B. Adelson, S. Dumais, & J. Olson (Eds.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systemspp. 439-440). New York: ACM.

Rada, R. (1995). Hypertext, multimedia and hypermedia. The New Review of Hypermedia and Multimedia, 1-19.

Rao, R., & Card, S. K. (1994). The Table Lens: Merging graphical and symbolic representations in an interactive focus+context visualization for tabular information. In B. Adelson, S. Dumais, & J. Olson (Eds.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systemspp. 318-322). New York: ACM.

Rao, R., Card, S. K., Jellinek, H. D., Mackinlay, J. D., & Robertson, G. D. (1992). The Information Grid, Proceedings of UIST'92 Conference of the ACM SIGGRAPH Symposium on User Interface Software and Technologypp. 23-32). New York: ACM.

Rao, R., Card, S. K., Johnson, W., Klotz, L., & Trigg, R. H. (1994). Protofoil: Storing and finding the information worker's paper documents in an electronic file cabinet. In B. Adelson, S. Dumais, & J. Olson (Eds.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systemspp. 180-185). New York: ACM.

Rheingold, H. (1991). Virtual reality. New York, NY: Summit Books.

Ribarsky, W., Bolter, J., Op den Bosch, A., & van Teylingen, R. (1994). Visualization and Analysis Using Virtual Reality. IEEE Computer Graphics & Applications, 14(1), 10-12.

Robertson, G. G., Mackinlay, J. D., & Card, S. K. (1991). Cone Trees: Animated 3D visualizations of hierarchical information. In S. P. Robertson, G. M. Olson, & J. S. Olson (Eds.), Proceedings of ACM CHI'91 Conference on Human Factors in Computing Systemspp. 189-194). New York: ACM.

Sarkar, M., & Brown, M. H. (1992). Graphical fisheye views of graphs. In P. Bauersfeld, J. Bennett, & G. Lynch (Eds.), Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systemspp. 83-91). New York: ACM.

Schaffer, D., Zuo, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S., & Roseman, M. (1996). Navigating hierarchically clustered networks through fisheye and full-zoom methods. ACM Transactions on Computer-Human Interaction, 3(2), 162-188.

Shneiderman, B. (1980). Software psychology: Human factors in computer and information systems. Boston: Little, Brown and Company.

Spoerri, A. (1993). InfoCrystal: A visual tool for information retrieval, Proceedings of IEEE Visualization '93pp. 150-157). Los Alamitos, CA: IEEE Computer Society Press.

Stone, M. C., Fishkin, K., & Bier, E. A. (1994). The movable filter as a user interface tool. In B. Adelson, S. Dumais, & J. Olson (Eds.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systemspp. 306-312). New York: ACM.

Strauss, S. (1995). Cybersickness: The side effects of virtual reality. Technology Review, 98(5), 14-16.

Stuart, R. (1996). The design of virtual environments. New York: McGraw-Hill, Inc.

Summitt, P. M., & Summitt, M. J. (1995). Creating cool 3D Web worlds with VRML. Foster City, CA: IDG Books Worldwide, Inc.

Thiel, U., & Muller, A. (1996). "Why was this item retrieved?": new ways to explore retrieval results. In M. Agosti & A. F. Smeaton (Eds.), Information Retrieval and Hypertextpp. 181-201). Boston: Kluwer Academic Publishers.

Thompson, R. H., & Croft, W. B. (1989). Support for browsing in an intelligent text retrieval system. International Journal of Man-Machine Studies, 30, 639-668.

Travers, M. (1989). A visual representation of knowledge structures, Proceedings of the ACM Hypertext '89 Conferencepp. 147-158). New York: ACM.

Turo, D. (1994). Hierarchical visualization with Treemaps: Making sense of pro basketball data. In B. Adelson, S. Dumais, & J. Olson (Eds.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems, Conference Companionpp. 441-442). New York: ACM.

Vacca, J. (1995, April). The Net's next big thing: Virtual reality. Byte, 28.

Wann, J., & Mon-Williams, M. (1996). What does virtual reality NEED?: human factors issues in the design of three-dimensional computer environments. International Journal of Human-Computer Studies, 44, 829-847.

Ware, C., Arthur, K., & Booth, K. S. (1993). Fish Tank Virtual Reality. In S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel, & T. White (Eds.), Proceedings of ACM INTERCHI'93 and CHI'93 Conference on Human Factors in Computing Systemspp. 37-42). New York: ACM.

Ware, C., & Franck, G. (1996). Evaluating Stereo and Motion Cues for Visualizing Information Nets in Three Dimensions. ACM Transactions on Graphics, 15(2), 121-140.

Waterworth, J. A., & Serra, L. (1994). VR management tools: Beyond spatial presence. In B. Adelson, S. Dumais, & J. Olson (Eds.), Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systemspp. 319-320). Boston: ACM Press.

Williamson, C., & Shneiderman, B. (1992). The Dynamic HomeFinder: Evaluating dynamic queries in a real-estate information exploration system, Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrievalpp. 338- 346). London: Springer-Verlag.