Hierarchical Data Visualization In Desktop Virtual Reality

by

David K. Modjeska

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto

Hierarchical Data Visualization in Desktop Virtual Reality
by David K. Modjeska

Ph.D. Thesis
Department of Computer Science
University of Toronto
2000

Abstract

While desktop virtual reality (VR) offers a way to visualize structure in large information sets, there have been relatively few empirical investigations of visualization designs in this domain. This thesis reports the development and testing of a series of prototype desktop VR worlds, which were designed to support navigation during information visualization and retrieval. Four methods were used for data collection: search task scoring, subjective questionnaires, navigational activity logging and analysis, and administration of tests for spatial and structure-learning ability. The combination of these research methods revealed significant effects of user abilities, information environment designs, and task learning.

The first of four studies compared three versions of a structured virtual landscape, finding significant differences in sense of presence, ease of use, and overall enjoyment; there was, however, no significant difference in performance among the three landscape versions. The second study found a hypertext interface to be superior to a VR interface for task performance, ease of use, and rated efficiency; nevertheless, the VR interface was rated as more enjoyable. The third study used a new layout algorithm; the resulting prototype was rated as easier to use and more efficient than the previous VR version. In the fourth study, a zoomable, map-like view of the newest VR prototype was developed. Experimental participants found the map-view superior to the 3D-view for task performance and rated efficiency.

Overall, this research did not find a performance advantage for using 3D versions of VR. In addition, the results of the fourth study found that people in the lowest quartile of spatial ability had significantly lower search performance (relative to the highest three quartiles) in a VR world. This finding suggests that individual differences for traits such as spatial ability may be important in determining the usability and acceptability of VR environments.

In addition to the experimental results summarized above, this thesis also developed and refined a methodology for investigating tasks, users, and software in 3D environments. This methodology included tests for spatial and structure-learning abilities, as well as logging and analysis of a user's navigational activities.

Acknowledgements

This research was made possible through funding from the Department of Computer Science at the University of Toronto. Please note that aspects of the present research have been published previously (Modjeska, 1998; Modjeska, 1999; Modjeska and Waterworth, 2000).

For design contributions to the research prototypes, I would like to thank Walter Frost and Russell Kroll, formerly of the Institute of Design at Umeå University. For first implementation of the research prototypes, I would like to thank Karel Moijson and David Steeman, formerly of the Hogeschool Gent.

For questions, advice, and support, I would like to thank the Interactive Media Lab in the Department of Mechanical and Industrial Engineering at the University of Toronto.

For encouraging my research interests and progress at a critical point, I would like to thank a former co-supervisor, Marilyn Mantei.

For contributions of expertise and perspective, I would like to thank the members of my thesis committee: Alberto Mendelzon, and Michiel van de Panne.

For supervision and inspiration, I would like to thank my guest supervisor, John Waterworth, in the Department of Informatics at Umeå University. The Informatics department as a whole contributed comments, companionship, and advice.

For weekly conversations about the meaning and implications of this research, I would like to thank my co-supervisor, monica schraefel.

For patience, supervision, and wit, I would like to thank my supervisor, Mark Chignell. His role in launching and completing this research, and developing opportunities for future work, has been central.

Finally, I’d like to thank my wife, Natalia Nygren Modjeska, for patience, inspiration, encouragement, and support during the final years of research.

Table of Contents

List of Tables

List of Figures

List of Appendices

"In cyberspace, information-intensive institutions and businesses have a form, identity, and working reality--in a word and quite literally, an architecture--that is counterpart and different to the form, identity, and working reality they have in the physical world. The ordinary physical reality of these institutions, businesses, etc., are seen as surface phenomena, as husks, their true energy coursing in architectures unseen except in cyberspace."

-- from "Cyberspace: Some Proposals" by Michael Benedikt

Chapter 1

Introduction

1.1 Research Context

As the volume and complexity of the information sphere grows, more advanced techniques are needed to visualize it. Such techniques, generally known as instances of information visualization, can be seen as moving information from abstractions to representations. Similarly, as the power and sophistication of computer technology increases, its ability to simulate the world also increases. Such approaches, generally known as virtual reality (VR), move information in the opposite direction, i.e., from concrete representation to abstraction. These two trends are starting to meet in the domain of information visualization using desktop VR. This meeting may eventually enable significant improvements in the everyday accessibility of large information sets.

Information visualization addresses the problem of representing various data types for the end-user, so that the data can more easily be understood, managed, and communicated. A variety of techniques has been developed to handle both structured and unstructured data, especially for scientific and industrial tasks. The hallmark of such techniques is a mapping from a source data domain to a destination visual domain, where task-dependent objects and relationships become apparent (visible) and available. A secondary goal is to offload part of the burden of conscious information processing to the human perceptual system (Robertson, Mackinlay, & Card, 1991). While often powerful, existing techniques generally place the user in a third-person perspective with regard to data, or perhaps in a first-person perspective in a limited visual space. Richer possibilities for exploiting human navigational knowledge are open for research and development. Moreover, it is unclear how some existing techniques would scale up for the size and complexity of very large data sets.

VR uses the techniques of computer graphics to present a model environment to the senses through a variety of computer media, ranging from full-body immersion to traditional desktop display. An informal definition of VR is as follows (Neill, 1994):

> Virtual Reality is a combination of hardware and software that allows you to see, move around in and manipulate computer graphics. . . . there are two basic components that a *true* virtual reality program *must* contain: 1) A first-person viewpoint that has complete movement at will in real time; 2) The ability to manipulate and/or change the virtual environment in real time (that is, the computer processes and displays the simulation as you experience it instead of playing pre-recorded sequences . . .).

The applications of VR to date have primarily been in engineering and architectural design (CAD), communication (virtual communities), and entertainment (computer games). All of these applications are based more or less directly on real-world phenomena, yet none of them explores in depth the potential of VR as a medium to represent abstract information. Recent research on wayfinding in VR suggests that human navigation skills transfer effectively from real into electronic worlds, in many cases (Darken and Sibert, 1996). If this is so, an opportunity exists to use VR for large-scale information visualization tasks. Such usage would have considerable benefits for accessing and communicating large information sets. While proposals and guidelines for such virtual worlds exist, few tested implementations have been described in the research literature (see Chen and Czerwinski, 1997, and Chen, 2000, for examples).

The best opportunities to use VR for large-scale visualization currently lies with browsing hierarchical data, for two reasons. First, the world's largest information structure is the Web. Individual Web sites often have a generally hierarchical structure (n.b. Lamping, Rao, & Pirolli, 1995). While software engines have greatly helped searching, browsing has received less support. Second, several well-known research prototypes of information landscapes for unstructured data exist. These prototypes are often constructed using statistical analysis techniques.

1.2 Research Questions

Within the context discussed above, a research project was initiated at the University of Toronto to investigate the effectiveness of desktop VR for visualizing hierarchical information. The emphasis was on user navigation, as this activity linked the problem domains of user psychology and system design. Some of the research was carried out in the Department of Informatics at Umeå University in Sweden, which has a research exchange program with the Department of Mechanical and Industrial Engineering at the University of Toronto.

During the course of this research, a number of themes were addressed, which might serve to orient the work presented later in the thesis. These themes can be presented as questions, as will be done here. Answers to some of these questions were found during the research project, and new questions naturally arose to suggest future research directions.

The first question would be the most general: *For which users is desktop VR an effective user interface (UI) for hierarchical information visualization?* To answer this question, a series of four user studies was conducted to test the effectiveness of prototype UI designs, as well as to identify significant dimensions of variation in individual users. The experimental tasks generally revolved around information searching, as this activity is important for both academic and industrial work. Related tasks assessed the learning of information structure and attitudes after exposure to experimental software designs. Details of these experiments will be presented later in the thesis (Chapters 3-6). Key dimensions of individual difference will be identified, especially in the last study. These are among the few task-focused studies published to date in the research literature on desktop or immersive VR.

Despite the novelty and attraction of a new interaction technique such as (desktop) VR, it is necessary to ask a related question: *How does searching in desktop VR compare with searching in other UI technologies?* Accordingly, the second study in this series focused on comparing the evolving research prototype with an established technique for hierarchy visualization – hypertext (Chapter 4). This baseline performance comparison resulted in a more robust research prototype for further study, as shown by the results of the fourth study (Chapter 6).

A particular dichotomy informed the beginning of this research, i.e., *what is an effective (and enjoyable) balance of text- and object-based representation of information structure in a virtual environment (VE)?* In other words, what is an appropriate balance of semantic and spatial structure? This topic has been researched in hypertext and cognitive science fields, but application to VR is new to the research literature. This situation offered an opportunity for the present research to consider the issue from software design and psychological perspectives. As mentioned, the research began by considering the question as a design issue. In the end, the research found significant elements of individual user difference on this issue. These issues will be discussed particularly in connection with the thesis' first and fourth studies (Chapter 6).

A further question, discovered in user studies, was this one: *To which experimental factors do users respond objectively vs. subjectively?* From the first experiment, it was clear that participants showed a distinct difference between attitude and behavior. That is, experimental conditions encouraging good task performance did not necessarily result in positive attitudes.

The preceding discussion raises another important research question: *What are the trade-offs between enjoyment and task performance?* Traditional HCI evaluation has focused on the latter, with its obvious benefits for measurable productivity. Yet the former has potential advantages for user motivation and focus. This issue persisted through most of the thesis research, with implications for both user behavioral evaluation and system design.

In a new design environment, a methodological question naturally arises: *What are appropriate strategies for user evaluation?* The HCI research literature contains substantial material about user evaluation methods in graphical user interface (GUI) environments. Unfortunately, not all of those methods apply to 3D environments, even on a 2D desktop (e.g., windowed event logging vs. 3D navigational logging). A research thread of the current project thus concerned the development of appropriate testing and evaluation methods for users of desktop VR. Related issues included task design and administration; identification and testing of individual differences; and recording and analysis of navigational activities. Working approaches to these issues were developed, as described in later chapters of the thesis (Chapters 3-6).

Returning to the computer side of HCI, another key question is this one: *What are effective structures and presentations of hierarchical information in desktop VR?* The design space is large, so criteria are needed for decision and evaluation. Presentation issues were easier to address, given a large body of research and practice in the field of visual design. Accordingly, trained designers assisted with presentation issues. Structural design was more problematic. An initial ad hoc solution proved relatively unsatisfactory for task performance, so a visualization algorithm from the research literature was implemented instead (Chapter 5). Implementation and testing of that algorithm were an important aspect of the current thesis research.

Implementation of the research prototypes proved a moderately substantial undertaking. Programming effort amounted to approximately seven full-time person-months of work. Development was undertaken in an object-oriented programming language for optimal maintenance, extensibility, and clarity of design. Traditional software engineering techniques were applied. The final software consisted of approximately 10,000 lines of code.

A final question concerns the attractiveness of novel, 3D information environments: *What is the value of the third dimension in a landscape metaphor for information visualization?* Computer games have shown the innate appeal of 3D environments for recreational purposes. Yet it is unclear that the engagement of games results in cognitive efficiencies that would benefit visualization. Once the thesis' research prototypes became sufficiently usable (as discussed above), the issue of the third dimension was considered explicitly. This issue will be discussed in connection with Study 4 (Chapter 6).

A research overview generated from the preceding questions is shown in Figure 1.1. The research focus in each of the project stages is shown in Table 1.1.

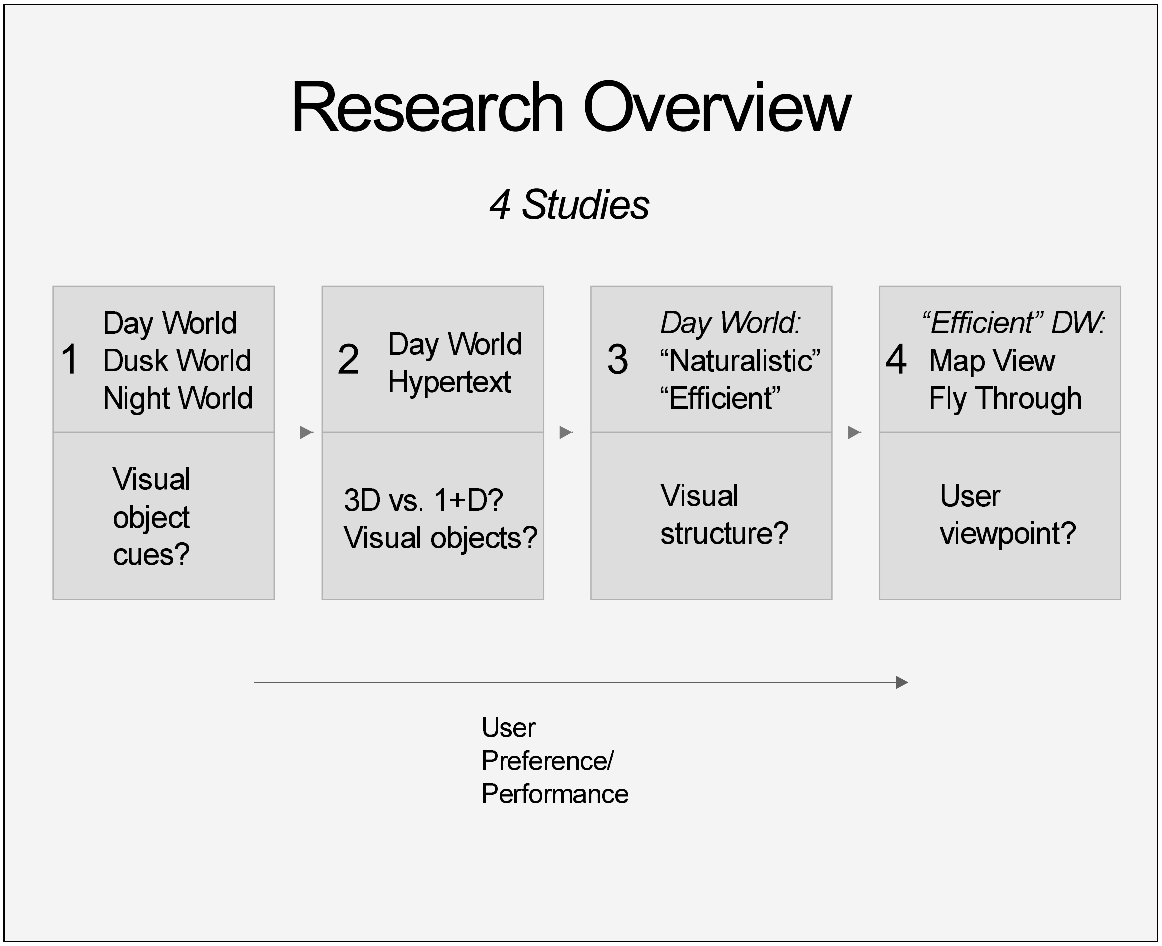

Figure 1.1. Thesis research overview.

Research Focus

		Study			
Area	**Issue**	**1**	**2**	**3**	**4**
Technology	Visualization design validation	X		X	X
Technology	Data-structure/UI relationships	X	X	X	X
Interaction	Search task methodology	X	X	X	X
User	Objective/subjective factors	X	X	X	X
User, Interaction	Spatial ability (with test and logs)				X
User, Interaction	Structure-learning ability (with test)		X	X	X

Table 1.1. Thesis research focus.

1.3 Thesis Perspective and Organization

In mediating between person and machine, the field of human-computer interaction (HCI) faces the implicit challenge of reconciling two sets of models, as well as finding a balanced perspective. In this thesis research, the question is whether to focus on a person operating on his or her environment, or on feedback offered by a computer system. The thesis will attempt to maintain a task focus, as previously mentioned. This focus should allow consideration of theoretical and practical issues on both sides, while striving towards solution to relevant problems. The emphasis will thus be placed on interaction style, which implicitly relates user and system.

The thesis will be organized as follows. First, related research will be reviewed, in order to situate the current project in a larger context. Second, four user studies will be described, each one including methodology, results, and discussion. Finally, general conclusions will be drawn about the current work and its implications for future research and applications.

This thesis research originated and was completed in the Interactive Media Lab of the Department of Mechanical and Industrial Engineering at the University of Toronto. The first three of the four user studies were performed in the Department of Informatics at Umeå University in Sweden.

Chapter 2

Related Research

2.1 Introduction

This review will begin with the fundamental topic of information structuring, and from there work towards an understanding of applications in information visualization.

In reviewing the research literature, the general goals of this chapter are as follows:

- To answer questions on the state of the art in information structuring, exploration, and visualization, and in particular, their areas of intersection
- To identify open research questions to be addressed
- To locate relevant work upon which to build
- To consider implications for research methodologies to be adopted

In general, this chapter will attempt to set boundaries for the thesis research, and to locate it in a broader context.

The remaining sections of the chapter will cover background issues, information structuring, physical and information exploration, information visualization, and general conclusions. The section on background issues will touch on some research in cognitive engineering and user-interface (UI) metaphors. These issues will appear in discussions of research in the body of this review.

1. The section on *information structuring* will discuss the properties and benefits of specific information structures, which affect both locomotion and wayfinding in electronic worlds. Second, it will suggest structural principles that could facilitate online navigation.

2. The section on physical and *information exploration* will introduce key research on navigation. It will describe general models, introduce the concept of mental maps, discuss wayfinding research, and review some theory and experiments in hypermedia.

3. The section on *information visualization* will discuss a variety of techniques and tools that could facilitate electronic navigation. It will then consider general principles that could inform the design of environments, tools, and interaction styles in electronic worlds. This section will discuss

research at XEROX PARC, review several hypermedia and spatial-world visualizations, and also consider the techniques of dynamic queries, semantic zooming, information landscapes, and multi-component (integrated) visualizations.

The chapter will conclude by summarizing key research issues and directions for future research.

2.2 Cognitive Engineering and UI Metaphors

Before discussing the main topics of this research review, it is worth noting briefly research in two areas affecting the fundamentals of human-computer interaction. The first of these areas is cognitive engineering, which links cognitive psychology with systems engineering. The second of these areas is UI metaphor, which links UI design with users' real experiences. These areas help to build a framework that bridges gaps between disparate research domains; this framework then supports the research discussed later in this review.

Cognitive engineering has two major goals: (1) to understand the principles of human action and performance that are relevant for developing engineering design principles, and (2) to design systems that are pleasant to use, as well as efficient, easy, and powerful. In cognitive engineering, according to Hutchins, Hollan, and Norman (1986), HCI directness (or distance) has two main aspects. First is the distance between a user's thoughts and the physical requirements of the information system. In any human-computer interaction, there is a distance between the user's goals and knowledge, and the level of description offered by the system: the gulf of execution (user → system) and the gulf of evaluation (system → user). Generally speaking, these gulfs are bridged by coordinating the system's interface with the user's mental structures and processes. Directness is inversely proportional to the cognitive size of these gulfs. Two types of directness are possible: semantic and articulatory. Semantic directness concerns the relationship between a users' intentions and the meanings of expressions in a UI "language" (interaction style). Articulatory directness concerns the relationship between the meanings of these expressions and their physical forms. The second main aspect of directness is the qualitative sense of engagement, that is, the sense of directly manipulating objects. With regard to engagement, direct engagement offers "a feeling of involvement directly with a world of objects rather than of communicating with an intermediary. The interactions are much like interacting with objects in the physical world. Actions apply to the objects, observations are made directly upon those objects, and the interface and the computer become invisible" (Hutchins et al., 1986). In summary, the most direct UI will minimize distance while maximizing engagement. From a design perspective, potential tradeoffs between distance and engagement can inform the development of information visualizations for particular tasks and users. In addition, the notions of distance and engagement are useful for interpreting participant responses to information exploration in experimental settings.

Although some researchers favor HCI interactions "[l]ike interacting with objects in the physical world" (Hutchins, 1995), virtual objects and interactions remain metaphorical. For this reason, it is worth considering the nature of metaphors. From a linguistic and philosophical point of view, Lakoff and Johnson (1980) claim, "our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature." They claim, "the essence of metaphor is understanding and experiencing one kind of thing in terms of another." A metaphor allows one concept to partially structure another, while at the same time concealing aspects of the structured concept. Metaphors don't establish complete mappings between concepts; rather, they enhance some aspects and suppress others. The most fundamental metaphors may emerge directly from physical, social, and cultural experiences. Each experience forms a coherent gestalt having several dimensions: participants, parts, stages, linear sequence, causation, and purpose. Often described by emergent metaphors, basic gestalts can structure more difficult concepts, e.g., emotions. Complex concepts often require two or more metaphors, each partially structuring the concept. That is, each subordinate metaphor represents an aspect of the complex concept, and these metaphors cohere together. Many metaphors fall into one of two major categories – orientational and ontological. Reflecting basic human experience in space, orientational metaphors use such oppositions as up-down, in-out, front-back, on-off, deep-shallow, and central-peripheral. (These metaphors can be used in online navigation.) Emotional states often rely on orientational metaphors for representation. Reflecting human experience with physical objects and substances, ontological metaphors allow people to define or bound undifferentiated direct experience. Such metaphors treat events, activities, emotions, and ideas as entities or substances. A personification is a special case of an entity metaphor.

2.3 Information Structuring

Having discussed some useful topics in perception and cognition, this section will review research on information structuring. The structure of an electronic world affects user navigation at both the articulatory and semantic levels. For locomotion (articulatory), information structure determines the routes that a user can follow. For wayfinding (semantic), information structure influences a user's conceptual model of a domain, as well as the perceptual cues that can be supported by the environment; conceptual models and perceptual cues strongly influence the navigational strategy chosen by a user. Another key issue is the nature of the information structure: a strong tension exists between semantic and physical structures. This section will review several topics in information structuring, with an eye to their impact on navigation: structure types and properties, hierarchical structures, semantic versus physical structure, experimental results, and designs for hypermedia and spatial worlds.

2.3.1 Semantic vs. Physical Models

In a human-computer interaction involving a user and an electronic world, there are perhaps three domains of structure: conceptual, semantic, and physical. The gap between conceptual and semantic structures, the “semantic distance” of Hutchins et al. (1986), will be discussed in connection with mental maps and wayfinding. The gap between semantic and physical structures (the “articulatory distance” of Hutchins et al.) will also be mentioned several times and merits further discussion. Several perspectives on this issue will now be reviewed.

Dillon, McKnight, and Richardson (1993) discuss the navigational roles of semantic intentions and physical representations:

> “Ultimately, we believe the idea of directly navigating semantic space has to be spurious. Semantic space is an abstract psycholinguistic concept which cannot be directly observed, only represented by way of alternative instantiations. By definition, semantic space is n-dimensional and practically unbounded. In order to visualize the semantic space it needs to be given physical representation and in so doing, it becomes at most three-dimensional . . . and physically bounded. . . . In effect we cannot navigate semantic space, at least not the way we navigate physical environments, we can only navigate the physical instantiations that we develop of the semantic space.”

The authors note that the preceding comments have two consequences. First, concepts of spatial navigation do apply meaningfully to physical representations of semantic intentions. Second, a well-designed document (or collection) should place its physical structure and its semantic contents in strong correspondence. Passini (1984) and Furnas (1997) made similar recommendations for wayfinding and effective view navigation, respectively.

In a related discussion, Kaplan and Moulthrop (1994) contrast physical and semantic “spaces.” Physical (“architectonic”) space, on the one hand, is regular, precise, stable, non-overlapping, and absolute. Familiar metaphors drawn from physics, architecture, and daily experience are derived from physical space. Semantic “space”, on the other hand, is connected to meaning, interpretation, and symbols; such space is unclear, ambiguous and unstable. It is created, for example, in the mind of a reader of hypertext documents. The authors claim that hypermedia systems generally favor architectonic spaces over semantic ones. Each type of space, however, implies the other one. In particular, navigation in an architectonic information space has semantic consequences, which are usually not supported by the system. Kaplan and Moulthrop set the design goal of generating “architectonic structures which, though still engaged in precise graphical mapping, are better adapted to the multiplicity of semantic space.” That is, the authors seek to display dynamic, local semantic information in hypermedia, in order to facilitate navigation.

Dillon et al. (1993) recommend designing a stable environmental structure with a close correspondence between semantic and physical structures. A large, well-structured Web site is an example of such an environment. The argument of Kaplan and Moulthrop (1994), by contrast, suggests designing dynamic environmental cues and tools to insert semantic information into physical structures. A large text collection with context-dependent, dynamic hypertext links (Golovchinsky, 1997) is an example of such an environment. These two sets of theoretical recommendations are complementary; empirical research would be beneficial for determining the relative importance and roles of the two approaches.

In related research, Shum (1990) discusses the properties and benefits of Euclidean (physical) and virtual (semantic) spaces. Virtual space, on the one hand, has only relative locations, since it has no global or fixed coordinate system. Euclidean space, on the other hand, has absolute locations. In any case, in electronic worlds of each type, an information node's location should relate meaningfully to its content. In general, spatial representation facilitates navigation and comprehension of structure. With 3D graphics, unfortunately, it can be difficult to represent meaningfully structural (semantic) relationships as relationships in space. Moreover, distance and directional cues for cognitive mapping can be lost in dynamic environments. Labeled dimensions can add semantics to spaces; but if more than three useful dimensions are available, multiple structures can undermine the stability required for cognitive mapping. In general, the author notes, the user’s task should determine the appropriateness of spatial imagery. Dillon et al. (1993) proposed that physical structures reflect semantic ones; Kaplan and Moulthrop (1994) proposed that dynamic, semantic information be represented locally in physical structures; and Shum (1990) argues for the nesting of physical and semantic structures. Again, further research is needed to clarify the relative roles of physical and semantic structures in electronic worlds, as will be pursued in this thesis.

2.3.2 Structure Types

In research based on Lakoff and Johnson (1980), Rennison and Strausfeld (1995) developed the Millennium Project, “a conceptual and computational approach for enabling understanding of a large, multidimensional set of information.” The system allows the user to explore a semantically structured space by navigation and manipulation. The project helps the user to understand relationships between items, seeking to create a dynamic virtual space with properties of mental ones. In particular, it explored the correspondences between metaphors of embodiment (up/down, forward/back, right/left, in/out) and informational structures. Conceptually speaking, the research is based on linguistic metaphor theory (Lakoff and Johnson, 1980) and cognitive science. The prototype uses a database of historical information, in which objects have attributes such as date, location, associations, cause-effect relationships, and size.

Adapting this work, we can identify five major types of information structure:

- *Set* structures sort items into distinct categories.
- *Graph* structures consist of nodes (vertices) and links (edges), arranged as a hierarchy or a network.
- *Relational* structures organize items by relative position.
- *Radial* (polar-geometric) structures organize items into center and periphery.
- *Linear* structures organize items on a continuum by a single attribute.

From these basic five structures, more complex structures can be constructed. Both homogeneous and heterogeneous combinations are possible. Cartesian planes and spaces, for example, are constructed with two and three orthogonal linear dimensions, respectively. A travel guide, a shopping mall directory, and the telephone yellow pages, as other examples, are often structured with a first-level linear (alphabetical) list of categories, each of which contains a second-level, linear (alphabetical) list of items. In information systems, one structure can often be implemented as a different one. The categories of Rennison and Strausfeld provide a useful conceptual vocabulary for designing and discussing navigational research. Lynch's (1960) five urban design elements (discussed in Section 4.3.1) influenced, and correspond to, that conceptual vocabulary.

Effective information structuring requires a selection of quantifiable attribute(s) from the data domain (Benedikt, 1991). He proposed four types of attribute: geographical, chronological, alphabetical, and domain-specific. Domain-specific attributes can be subdivided in two ways: arranged into categories, or along a continuum (Wurman, 1989). Rennison and Strausfeld (1995) suggest a natural correspondence between domain attributes and types of information structure. This correspondence can be useful in the design of electronic worlds, especially for establishing a strong correspondence between semantic and physical structures (Dillon et al., 1993).

Regardless of the environmental structure or user strategy, navigation can be difficult in complex, large-scale information structures. Furnas (1997) studied this problem and developed some design recommendations. He first defines two activities, which are essential for good user navigation:

- *view traversal* - an iterative process of viewing information, selecting a seen item, and moving to it, in order to follow a path through the information structure (locomotion)
- *view navigation* - view traversal where selections are informed and reasonable for target-seeking (wayfinding)

For this discussion, assume a structure characterized by a logical structure graph; a view graph reflects the portion of the overall structure visible from a particular point. Effective view traversal, then, requires a small view and short paths (small diameter) in the view graph. During design, inefficient view-traversable structures may be improved by adding a traversable infrastructure (c.f., Card, Robertson, & Mackinlay, 1991; Dieberger, 1995; Waterworth, 1996). A target's information residue or scent, per information foraging theory (Pirolli and Card, 1995), is defined as the remote indication of a target in outgoing link information throughout an information structure. View navigation then requires that outgoing link information be everywhere well-matched; that every node have good residue at every other node; and that outgoing link information be small. These requirements for view navigation imply the representation of many target sets, as well as an interlocking web of set representations, in which residue is a shared resource. High-level semantics thus play an important role in navigable structures, perhaps being used to subdivide an information structure efficiently.

Furnas critiques a few information navigation schemes:

- bad - WWW (bad residue, bad diameter); simple scrolling lists (bad diameter)
- mixed - geometric zoom (good diameter, bad residue)
- good - semantic zoom (better residue), 3D (better diameter), fish-eye views (shorter paths), balanced rooted trees (short paths and maybe simple semantics)

Although theoretical and focused on discrete information environments, Furnas' work is compatible with Passini's behavioral work in continuous physical environments. First, Furnas suggests the importance of high-level semantics in structuring an environment, for the purpose of facilitating a user's cognitive mapping. Second, Furnas notes the importance of good information residue at every point, which would provide the perceptual cues needed for good navigation planning. Third, Furnas mentions the importance of short paths, so that navigation plan execution is feasible.

2.3.3 Hierarchical structures

In research on navigation, information structuring, and information visualization, hierarchical structures feature prominently. According to Furnas and Zacks (1994), useful hierarchical structures fall on a continuum between trees and general graphs. Trees are conveniently planar, easy to traverse, and semantically analogous to the useful processes of abstraction and aggregation; unfortunately, trees support only fixed navigation and organization. At the other end of the scale, general graphs allow flexible navigation and multiple organizations; unfortunately, graphs are difficult to lay out, easily cause user disorientation, and are hard for users to abstract well. Intermediate structures, DAGs naturally support semantic abstraction; like general graphs, though, DAGs suffer from a lack of constraints. (Another intermediate structure is the pre-tree, Mukherjea, Foley, & Hudson, 1995.) Furnas and Zacks have developed a hierarchical structure between a tree and a DAG, called a multi-tree. This structure can be naturally interpreted as showing multiple contexts; it also supports the re-use of hierarchical components. Simply described, a multi-tree is not a tree, but the descendants of each node form a tree. The ancestors of each node form an inverted tree. Looking downward from a multi-tree node, one sees semantic content; looking upwards, one sees context. The multi-tree structure facilitates browsing and representation (including fisheye views), but it is difficult to display fully. In light of Furnas' work on effective view navigation, the multi-tree seems a good structure for complex, large-scale information environments. The multi-tree could support global information residue through multiple sets of target representations with hierarchical semantics (looking down the graph), as well as complete anchoring context (looking up the graph).

Additional support for intermediate hierarchical structure comes from architectural theory by Alexander (1965), who argues strongly for viewing and designing cities as "semilattices" (DAGs), rather than as trees. The graphical nodes in this argument are urban design elements, from the macroscopic scale of regional planning to the microscopic scale of interior design. Alexander claims that traditional cities are organized as DAGs; artificial (planned) cities are trees. Alexander complains that planned cities lack the vitality and aesthetic appeal of older cities. He notes that "the idea of overlap, ambiguity, multiplicity of aspect and the semilattice are not less orderly than the rigid tree, but more so. They represent a thicker, tougher, more subtle and more complex view of structure ." DAGs are more difficult to visualize than are trees; for this reason, Alexander feels, both designers and residents tend to conceive of cities as trees. As discussed above, research by Furnas on multi-trees and effective view navigation supports the adaptation of Alexander's principles for hypermedia design.

Alexander (1979) later extended the ideas about hierarchical urban structure to a general philosophy of design. These ideas have been influential in urban design and software engineering. In Alexander's view, well-designed cities and buildings manifest a DAG structure of design elements. These elements are recurring patterns of events in time, interlocked with geometric patterns in space. That is, the patterns are four-dimensional. According to Alexander, these patterns constitute the fundamental structure of the built environments. Each such pattern has three parts: a context, a system of forces (a problem), and a configuration (a solution). In practice, patterns behave like genes that govern design, and like cultural languages for building. Good patterns are effective and enhance human activity. Through research and practice, designers can discover, test, improve, and validate the patterns. A vital pattern language is shared by a community of people, both informing and interpreting their environment. Such a language ideally has a rich and complex structure, which evolves with changing needs. As well as supporting rich structure, Alexander's "pattern language" offers a useful framework for analysis and design. It also underscores the importance of several design values: multi-dimensional (spatio-temporal) thinking, evolutionary thinking, and social awareness.

2.3.4 Experimental Results

Two behavioral studies offer an empirical basis for some of the theoretical claims discussed above.

Valdez, Chignell, and Glenn (1988) investigated the role of landmarks in user orientation during hypermedia browsing. The authors note the advantages of fisheye lens models (Furnas, 1986) for balancing flexible information access and continuous user orientation. Landmarks are particularly useful for fisheye models of hypermedia structure. Accordingly, Valdez et al. conducted two user experiments that involved recall tasks, cognitive structure mapping, and spatial path construction. The authors found that "connectivity and levels of abstraction may be the best predictors of landmark quality." Landmarks were even important to users who lacked strongly spatial cognitive structures. When users lack strong cognitive maps, the authors speculate, fisheye views may require a semantic basis, rather than a physical one based on hypermedia link structure. General informational landmarks could represent concepts such as categories and icon symbols. Accordingly, the authors suggest that future research might enhance identified landmarks with visual semantics, to make them easier to interpret and use.

The importance of landmarks for physical navigation was noted by Lynch (1960) and Passini (1960); Valdez et al. confirm this finding for hypermedia. Semantically-enhanced landmarks would follow the recommendations of McKnight, Dillon, and Richardson (1991) and Kaplan and Moulthrop (1994). Accordingly, semantic enhancement appears to be a fruitful direction for research on hypermedia route maps. Survey map knowledge may always remain problematic for users, on account of the restricted two-dimensionality of hypermedia structures. Meanwhile, the authors' use of hypermedia connectivity to identify landmarks has influenced subsequent research on Web structure navigation and visualization, which will be discussed below.

On the basis of the research by Valdez et al., GIT researchers developed an automated procedure to determine whether a Web node is a structural landmark (Mukherjea and Foley, 1995). The procedure considers the number of other nodes reachable via directional links in one or two steps. Four measures are defined:

- *outdegree* (O) - one step forward
- *indegree* (I) - one step backward
- *second-order connectedness* (SOC) - two steps forward
- *back second-order connectedness* (BSOC) - two steps backward

The procedure has two steps:

1. Node Importance = ((O + I) × Weight1) + ((SOC + BSOC) × Weight2), where Weight1 + Weight2 = 1. The researchers had the best results with Weight1 = 0.4 and Weight2 = 0.6
2. The given node is a landmark iff Importance > 10 % of the total number of nodes.

2.3.5 Designs for Hypermedia and Spatial Worlds

Having reviewed some theoretical issues and experimental results about information structuring, it is worth stepping back to consider these structures in a context of design and use. This consideration provides a sense of potential information domains and user tasks, as well as a sense of some key design issues. This section will discuss a historical proposal for global hypermedia, a brief overview of the WWW, and a general proposal for spatial information worlds.

Many people have contributed to the development of hypermedia, but Bush (1945) is generally credited with proposing the general concept. In a seminal article, he suggested a device called the “memex”:

> Consider a future device for individual use, which is a sort of mechanized private file and library. It needs a name, and to coin one at random, ``memex" will do. A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.

This proposed device is sometimes seen as a design sketch for the World Wide Web. Bush proposed the memex as a mechanism to allow people -- especially scientists -- fast and associative access to the huge store of accumulated human knowledge. Such access, he felt, is required for humanity to incorporate historical knowledge for solving present problems, to automate the management of a complex civilization, and to supplement limited human memory with an external storage medium. Bush envisioned memex users as both writers and readers, creating and following informational paths. (Note the navigational metaphor.) The author considers the creation of the hypermedia link to be the key innovation, since he felt that the human mind operates by association, rather than by indexing (as in databases). The technical details of Bush’s proposal are now dated, so they will not be reviewed here.

An implementation of a “memex,” the World-Wide Web was originally designed to support shared knowledge for Internet-based collaboration at CERN, the European Particle Physics Laboratory in Geneva (Berners-Lee, Cailliau, Luotonen, Nielsen, & Secret, 1994). The WWW includes the idea of an information world, an address system (URL), a network protocol (HTTP), a markup language (HTML), and the data on the Internet. Designed to be flexible and extensible, the Web is undergoing exponential growth. According to estimates, the Web had more than a billion documents and more than 100 million regular users in

2000. Current uses involve a wide range of work and leisure activities around the globe. Anecdotally, the Web is a general graph in structure, and individual sites are often hierarchical (n.b. Lamping et al., 1995). The Web takes advantage of technological innovations in computing that weren't foreseen for the memex. Other differences include the Web's incorporation of multiple media other than text, as well as the Web's increasing commercial emphasis.

Complementing these proposals for global hypermedia, a theoretical model for global spatial worlds (cyberspace) has been proposed by Benedikt (1991). His research draws on physics, mathematics, psychology, and architecture. Benedikt describes ideal cyberspace as follows:

> Cyberspace is a globally networked, computer-sustained, computer-accessed, and computer-generated, multidimensional, artificial, or "virtual" reality. . . . This information derives in part from the operations of the natural, physical world, but for the most part, it derives from the immense traffic of information that constitute human enterprise in science, art, business, and culture. . . . The dimensions, axes, and coordinates of cyberspace are thus not necessarily the familiar ones of our natural, gravitational environment: though mirroring our expectations of natural spaces and places, they have dimensions impressed with informational value appropriate for optimal orientation and navigation in the data accessed.

Benedikt proposes a series of seven principles that adapt key aspects of the physical world for information worlds. He seeks to find reasonable tradeoffs between useful semantic structures and familiar physical ones. The principles cover five topological attributes of virtual space: dimensionality, continuity, curvature, density, and limits. The seven principles are:

1. *Exclusion* - no two objects can be in the same place at the same time.
2. *Maximal Exclusion* - from n-dimensional data, choose as extrinsic dimensions (virtual space and time) those that will minimize the violations of the principle of exclusion.
3. *Indifference* - the felt realness of the environment depends partly on its indifference to the presence and actions of users, i.e., the environment's persistence and autonomy. (User-friendliness is a related but not identical design consideration.)
4. *Scale* - the maximum velocity of user motion in cyberspace is an inverse, monotonic function of the complexity of the world visible to him.
5. *Transit* - in general, travel between two cyberspace locations should occur phenomenally through all intermediate points, and at a cost proportional to a distance measure.
6. *Personal Visibility* - users should always be visible to other nearby ones, and users should have control over the extent to which other users are visible to them.
7. *Commonality* - virtual places should be objective for a community of users.

In general, Benedikt suggests that a useful cyberspace will blend the "fine grain and powerful monotonic ordering of natural space dimensions" (physical structures) with the "pragmatic groupings of information classes, partially ordered, of structures" (semantic structures). Similarly, Benedikt observes that electronic worlds contain both navigation and destination data. Although GUIs separate such data, cyberspace "offers a deep, spatially continuous environment rich enough for objects to be ambiguously navigational and 'destinational'--switching, phenomenally from one to the other as a function of user proximity, motivation, and attention, quite like reality." Like Alexander (1965) and Furnas (1997), Benedikt supports the use of multi-purpose design elements.

Benedikt envisions an information world with substantially the same activities as those noted for the hypermedia above. The form of the proposed world, however, is quite different. Benedikt emphasizes physical models because of innate human skills in dealing with the physical world. He also recognizes the need to integrate semantic information for navigation, as other researchers have noted (Kaplan and Moulthrop, 1994; McKnight et al., 1991; Shum, 1990). Benedikt adopts a novel approach, proposing a sort of virtual physics partly governed by information semantics. That is, the phenomenal (experiential) character of a world should reflect the information content in its dimensional axes, object positions, possible velocities, relative distances, entity visibility, etc. This approach to semantic and physical structures is more spatially sophisticated than some other approaches (Kaplan and Moulthrop, 1994; McKnight et al., 1991; Shum, 1990), but it is probably more computationally demanding as well.

2.3.6 Discussion

This section has reviewed some research on information structuring from a variety of disciplines - architecture, hypermedia, computer science, and psychology. Among this variety of approaches, some common concerns and conclusions have arisen:

- Structural issues affect both locomotion and wayfinding, which have different requirements for effective navigation.
- Several principles from urban design are applicable to electronic worlds:
  - the use of fundamental structures (graphical, linear, radial, set-like, and relational)
  - the value of hierarchical structures that are intermediate between trees and general graphs
  - the importance of using domain semantics to support the user's cognitive mapping
  - the importance of good information scent or residue
  - the importance of designing environments to support the execution of navigation plans
  - the value of multi-purpose design elements

- A strong tension exists between physical and semantic structures. Several approaches have been proposed - correspondence, parallel, nesting, and interpenetration. Researchers agree, though, on the importance of designs that incorporate both types of structure.
- Route maps are useful in hypermedia worlds, but survey knowledge is problematic. Semantic knowledge may be more useful than survey knowledge in this domain.
- Relatively little empirical research has considered key issues of information structuring in electronic worlds. Further research is required in this area.

2.4 Physical and Information Exploration

The previous section of this chapter discussed ways to structure and constitute an electronic world. This section will review research on navigation and information-seeking behavior, both in the physical world and in electronic environments. In many ways, the literature discussed here is the most important in the review, especially the material on wayfinding. The section is divided into five parts: general models for information exploration, mental maps, wayfinding, hypermedia environments, and spatial ability.

2.4.1 General Models

For information systems, Waterworth and Chignell (1991) propose a concise, general model of user-initiated exploration. The model has three orthogonal dimensions, which define a cubic design space:

1. *Structural responsibility* concerns navigation versus computer-mediated information retrieval (IR). The issue is user versus system responsibility for organizing information during search. Navigation is traditionally associated with hypertext, while computer-mediated information retrieval is traditionally called IR.

2. *Target orientation* concerns querying versus browsing. The issue is the specificity of the information requirement in the user's mind. If the user has a definite target, the activity is called querying; if the user has a less specific target and some interest in serendipity, the activity is called browsing.

3. *Interaction method* concerns descriptive specification versus referential specification. The issue is the user's interaction style with the system to obtain information. In descriptive specification, the user describes the target in text, speech, a database query language, etc. The method is primarily linguistic and is often associated with a conversational metaphor of HCI (e.g., a command-line UI). In referential specification, the user interacts with the system non-descriptively, by pointing,

clicking, selecting, or indicating a choice at each interaction stage. This method is primarily indicative and is often associated with a direct manipulation metaphor of HCI (e.g., a graphical user interface or GUI).

In designing an interface, the authors note the importance of finding appropriate combinations of the three exploration dimensions for a specific application and a user population. Intermediate points along the three axes are useful and common. In considering navigation in electronic worlds, this review focuses primarily on referential, browsing, navigational UIs, which are located at or near one corner of the design space proposed by Waterworth and Chignell (1991). While useful, the proposed model could be improved by taking greater cognizance of the user's understanding of the information system (Golovchinsky, 1997).

For evaluating the cost structure of such information navigation and gathering, Pirolli and Card (1995) proposed information foraging theory as a way to analyze human activities with information access technologies. Optimal foraging theory in biology and anthropology analyzes the adaptive value of food-foraging strategies. The information version analyzes tradeoffs in the value of information gained against activity costs in HCI, i.e., information cost structure. The theory considers a time scale from 100 ms. (cognition) to several months (interpersonal activity). Both the external information environment and human adaptation are considered. Optimality models generally include three types of assumption – which problems to analyze (decision), how to evaluate choices (currency), and the relationship between decision and currency assumptions (constraints). The information foraging task is usually embedded in another task context, which determines the value and cost structure. Users rely on representations of content as an "*information scent* whose trail leads to information of interest" (Pirolli, 1997). In general, human behavior exhibits bounded rationality and satisficing. Optimality models thus describe the possibility of a niche, "a possible advantageous adaptation if not blocked by other forces" (Pirolli and Card, 1995). As described by these researchers, such models include the information patch model, the information diet model, and the dynamic foraging model. Pirolli and Card performed a study using XEROX PARC's Scatter/Gather document browser with the large NIST Tipster document corpus. Results showed that user behavior matched simulation based on optimality models: compared to standard measures, Scatter/Gather user gains and HCI time-cost tradeoffs correlated well with a dynamic programming version of an information foraging model. A comprehensive summary by Pirolli and Card (1999) is available.

Information scent was used in a prototype system to analyze and predict user behavior and Web site usability (Chi et al., 2000). For this purpose, the system incorporated a visualization of a Web site (www.xerox.com) and its usage patterns; predictive modeling of user flow by information scent and agent-based simulation; and usability metrics based on document sequences in navigation paths. Although limited, the prototype has demonstrated the power of combining such methods for Web (and potentially VR) researchers and designers.

Information foraging theory has three implications for navigation in electronic worlds. First, it suggests that information cost structure can be productively considered in the design of online environments. Second, an electronic world (or sub-world) might be designed to support a foraging optimality model, depending on intended users and tasks. Third, the need for good information scent underscores the importance of appropriate environmental cues. As an adaptive, information-processing activity, information foraging is compatible with wayfinding as presented below (esp. Passini, 1984).

2.4.2 Mental Maps

During physical and information exploration, mental or cognitive maps are essential for human storage and use of environmental information. Such maps are the dominant concept for internal representations of large-scale environments.

Research by Stevens & Coupe (1973) provided evidence of hierarchical representations of environmental knowledge. Their experiment showed that human directional judgments about geographic locations can be substantially distorted by large, surrounding regions. The experimental tasks were all directional, asking subjects to indicate the relative direction of one location from another. Location pairs included San Diego and Reno, Nevada; Portland, Oregon and Toronto; and Montreal and Seattle. As a control condition, the experiment was repeated under laboratory conditions with fictitious cities and countries. In all cases, people consistently and significantly biased directional judgments towards the containing geographic region. These errors appear to arise from hierarchical mental representations of spatial information. The errors occur because spatial relations are often inferred from other facts. In the experiment, relative directions between cities were inferred from relative directions between states and countries. By way of explanation, the authors note that it is not economical to store all possible spatial relationships in long-term memory: an expanding knowledge base would risk exponential expansion. Hierarchical representation, therefore, supports efficiencies of storage and searching.

Further evidence for hierarchical representations of spatial knowledge was provided by Chase (1983). His study of taxi drivers in Pittsburgh reported several distortions in spatial judgments. The study investigated whether experts, after years of experience, developed more accurate cognitive maps than did novices. In laboratory tasks, drivers sketched a quadrilateral set of streets; this task generated rectilinear simplifications. Subjects also listed and drew neighborhoods, which were consistently clustered. Subjects also estimated Euclidean distances, which were amplified across neighborhood boundaries. Finally, subjects placed neighborhoods on a map, which generally skewed them towards a downtown landmark. In a task outside the laboratory, drivers were asked to point towards downtown; but they usually could only indicate the nearest main access street. In and outside the lab, drivers generated or improved routes. In this task, experts generated shorter routes and detours, as well as demonstrating superior picture recognition and street naming. (These experts may have exhibited superior *survey knowledge*, which will be discussed in connection with Thorndyke and Hayes-Roth below.) Both experts and novices performed better outside the laboratory. These results argue against so-called “maps in the head”; rather, they indicate that *large-scale environmental representations have a hierarchical organization.* Expertise seems to consist of an expanded knowledge of neighborhoods, streets, and environmental perceptual cues.

The basic concept of mental maps was extended by Downs and Stea (1973) to encompass definition, development, and use. Their definition of cognitive mapping is as follows:

> Cognitive mapping is a process composed of a series of psychological transformations by which an individual acquires, codes, stores, recalls, and decodes information about the relative locations and attributes of phenomena in his everyday spatial environment . . .

In this context, a cognitive map represents a functional analogue of a cartographic map. This view is compatible with the propositional view of mental map representation. Cognitive versions probably use a variety of signatures, which are sets of encoding and decoding operations. These signatures resemble cartographic maps, linguistic signatures, and visual imagery. While the human spatial environment is large, complex, and dynamic, human information-processing capabilities are limited. For spatial behavior, therefore, people develop cognitive maps containing two basic types of information about environmental phenomena:

- locational information (distance and direction)
- attributive information (description and evaluation)

Locational information, on the one hand, creates a subjective spatial geometry. The first locational component, distance, can be measured in many ways, since humans are sensitive to cost-benefit issues. The second component, direction, is used less consciously. People must often translate between polar spatial information and Cartesian cartographic maps. Attributive information, on the other hand, describes the nature of phenomena. Such attributes are of two types: descriptive (affectively neutral) and evaluative (affectively charged). Imperfect as knowledge containers in general, "cognitive maps are complex, highly selective, abstract, generalized representations in various forms. . . . incomplete, distorted, schematized, and augmented, and we find that both group similarities and idiosyncratic individual differences exist" (Downs and Stea, 1973). Cognitive maps are utilitarian. In practice, people acquire such maps by integrating information from all sensory modalities, and from direct, vicarious, and inferential sources. Cognitive maps tend to evolve over time.

The ideas of Downs and Stea have two implications for navigation in electronic worlds. First, users may have difficulty developing mental maps of very dynamic electronic worlds. Research is needed to establish the limits and mechanisms of mental mapping in such systems. Second, users may be highly skilled in developing and using attributive information in mental maps. If so, such information should be incorporated into electronic worlds to assist with navigation.

Experiments by Thorndyke and Hayes-Roth (1983) show differences in spatial knowledge acquired from maps and physical navigation. The authors propose a broad distinction between two types of spatial knowledge, which we will preface with a third type:

- basic feature recognition or *landmark knowledge* (McKnight et al., 1991)
- procedural descriptions of *route knowledge*
- *survey knowledge* of an environment's topographic properties

In opposition to the propositional view previously discussed, Thorndyke and Hayes-Roth assume an isomorphism (formal correspondence) between a physical map and its mental representation. The authors propose that physical and mental map users, on the one hand, assume a perspective above the horizontal domain. Visual-type search identifies absolute and relative object locations, while visual-type measurement allows users to assess Euclidean distances and relative directions. Navigators, on the other hand, have a perspective in the domain. Mental simulation of navigation permits users to identify route distances, and mental computational procedures permit users to assess Euclidean distances and relative directions. With increased navigation experience, users reorganize route knowledge into survey knowledge, although from a perspective in the plane.

To investigate these proposals, Thorndyke and Hayes-Roth (1982) conducted experiments with 48 subjects in the large headquarters of the Rand Corporation. Experimental tasks included placing features on an external map; pointing towards unseen locations; and estimating Euclidean and route distances. Two conditions pertained - map learning and navigational learning. The study derived detailed procedures for users' spatial judgments, which included distance estimates, orientation, and object location. Results showed that map learners judged location more accurately than orientation, because of the perspective shift required for them to judge orientation. Similarly, navigation learners judged orientation more accurately than location. Furthermore, with minimal learning, navigation learners judged orientation more accurately than did map learners. With minimal learning, map learners judged location more accurately than navigation learners, but experience removed this difference. Clearly, in the short term, maps offer benefits for perceiving and learning global relationships; with much experience, however, navigation generates comparable survey knowledge and offers the benefit of superior cognitive maps. The authors suggest further research to account for variations in environment uniformity and individual spatial ability.

The work of Thorndyke and Hayes-Roth has several implications for navigation in electronic worlds. First, the task and environment will determine the type of spatial knowledge required by the user. The system must offer appropriate support for this knowledge in the form of perceptual cues, environmental structure, tools, and/or social resources. Second, for those tasks and environments where survey knowledge is required, it can be acquired either by user experience or referencing external maps. For novice users, external maps will be required. Third, certain mental operations seem to treat mental maps functionally as if they were cartographic. External maps may facilitate these operations, as well as helping to overcome certain user errors (Stevens and Coupe, 1973; Chase, 1983).

In navigation research described above and below, a common research technique is “cognitive cartography” -- map sketching from memory to explore cognitive systems. Research by Canter (1977) considers two key questions. First, do maps model aspects of cognition? Second, what is implied by mapping procedures for exploring cognition? Canter describes a map as "an efficient means of recording any explicit spatial distribution of phenomena and their attributes .” Maps have two forms, conveying either route or survey knowledge. First, a *plan* (route) is a sequence of locations through space. Second, an *account* (survey) describes the overall arrangement of places (and attributes) in space. According to Canter, four processes are needed to transform spatial information into an external map: (1) orientation, (2) miniaturization, (3) projection, and (4) symbolization. These operations abstract experience into a form that supports action. By extrapolation, internal representations and subsequent sketch maps are also transformations of experience. As suggested above, sketch maps offer much potential for distortion. Forms are usually simplified, as part of structuring experience. In general, sketch maps show two types of information, which seem to be present in human cognitive systems: the links between places (route knowledge), and place locations (survey knowledge). Large differences in individual spatial ability are reflected in the accuracy and complexity of sketch maps. Because of these differences and the tendency to simplify mental maps, common external maps (e.g., in subway systems) use simplifying systems. Canter concludes that the psychological processes for dealing with place are cognitive, rather than perceptual. Accordingly, mental maps are not visual records, rather conceptualizations or frameworks that function as maps. This view is compatible with the propositional view of mental map representations that was discussed previously. Canter’s work suggests that electronic worlds might facilitate the creation of mental maps by explicitly supporting the four stages of spatial transformation: orientation, miniaturization, projection, and symbolization. His work also recommends the simplification of online maps for novice or spatially unskilled users, or at least the progressive disclosure of spatial complexity. For an empirical study of mental mapping in user-structured electronic worlds, rather than system- or designer-structured worlds, see Czerwinski, Dantzich, Robertson, & Hoffman (1999).

2.4.3 Wayfinding

2.4.3.1 Physical World

Partly for improving individuals' mental maps, urban planner Lynch (1960) researched psychologically important, structural patterns that affect complex information environments. This research has influenced the fields of urban design, psychology, and computer science (information visualization). Lynch proposes that a city's visual legibility or "imageability" strongly affects the satisfaction and navigational effectiveness of its residents. He argues that architects and planners can improve urban environments by following specific guidelines for legibility. In effect, he is proposing *structural design principles to improve users' mental maps*. An environmental image possesses three major components: identity, structure, and meaning. *Imageability* may then be defined as follows:

> . . . that quality in a physical object which gives it a high probability of evoking a strong image in any given observer. It is that shape, color, or arrangement which facilitates the making of vividly identified, powerfully structured, highly useful mental images of the environment."

To investigate the images held by city residents, Lynch conducted research in three U.S. cities. This research involved resident interviews, map sketches (cognitive cartography), trip and feature `descriptions, and photo classification; trained observers also conducted field analysis of the cities. These disparate data were assembled into cartographic maps. Lynch's methodology is notable for assembling heterogeneous research signatures (sets of cognitive encoding and decoding operations) into unified cartographic maps. The results indicate five key structural components of imageability: paths, edges, districts, nodes, and landmarks:

- *Paths* are an observer's channels of movement, and are the most important city elements (e.g., roads).
- *Edges*, other linear elements, act as boundaries between two areas (e.g., rivers).
- *Districts* are two-dimensional city sections into which observers mentally enter, and which have distinctive characters (e.g., neighborhoods).
- *Nodes*, smaller in scale, are strategic city spots into which observers can enter, and which act as the focus of transportation lines or other urban characteristics (e.g., plazas).
- *Landmarks* are single objects external to viewers, and upon which longtime city residents rely strongly (e.g., towers).

As previously noted, these five urban elements (Lynch, 1960) correspond to the information structure types of Rennison and Strausfeld (1995): path element / graph structure, edge / line, district element / set structure, node element / radial structure, and landmark element / relational structure.

Lynch notes that a satisfying city form should weave these five image elements into a strong pattern. The resulting overall image typically comprises a set of smaller images that are organized hierarchically by scale. This view strikes a compromise between the propositional and analog mental map representations discussed above. Among different residents, a city's image varies in quality with regard to density of detail, concreteness, and structural precision. Lynch offers specific guidelines for each of the five key urban elements, as well as identifying formal qualities upon which a designer may operate to create image elements of the recommended types. To create an overall form of clearly related parts, the author suggests three techniques: composing the region as a static hierarchy, relating smaller things to one or two dominant things, or organizing the region as sequence or temporal pattern.

Being relatively formal in nature, Lynch's structural recommendations for the physical world may improve wayfinding in electronic worlds as well. Research is needed to validate this claim (e.g., Chalmers, Ingram, & Pfranger, 1996; Ingram and Benford, 1995). In this vein, Shum (1990) notes that users' cognitive maps of hypermedia clearly include landmarks, paths, and nodes; districts and edges could be usefully added via UI metaphors. Finally, Lynch's research methodology could be adapted and extended for electronic worlds.

Later research by Passini (1984) built on Lynch's wayfinding research. Where Lynch focused on the physical and spatial characteristics of the urban environment, Passini focused on human information processing, environmental meaning (functional, sociocultural, and sensory), and information design. Like Lynch's, Passini's work was a mix of theory and experiment that generated design guidelines. Passini sought to describe the wayfinding mechanism and to explore the information processing that relates a person to his environment. Passini conducted several experiments with subjects in complex urban centers of downtown Montreal. In Passini's definition wayfinding has three iterative components:

- cognitive mapping or information generation to understand the environment
- decision making to structure and plan actions
- decision execution to transform decisions to behavioral actions

In general, a person's wayfinding is either linear, using directional signs, or spatial, using organized information about the complete setting. Passini found that individuals have a preference for either route or survey information; this finding extends the research of Thorndyke and Hayes-Roth (1982). Depending on individual cognitive bias, survey thinkers can use route information, but not vice versa. To facilitate wayfinding in general, designers should locate necessary information so as to be available during decision planning. Appropriate signs, specifically, can facilitate route-style wayfinding.

Of Lynch's five city elements (paths, landmarks, districts, nodes, and edges), Passini found that wayfinding requires mostly paths and landmarks. Passini proposes three elements to structure Lynch's elements into a coherent whole:

- An *organizational principle* fosters imageability and provides cognitive economy.
- *Spatial enclosure* enhances memorability, and supports inferences about content and structure.
- *Spatial correspondence* between environment features helps in developing integrated cognitive maps.

Passini found external maps to be useful for wayfinding. He suggests that they should support both route and survey knowledge. For developing route knowledge, a map should support three operations: cardinal or relative orientation of the map to the setting, locating the destination, and developing and memorizing a route. For developing survey knowledge, the map should feature Lynch's five elements and Passini's three elements.

Because it focuses on human information processing and information design, Passini's research may be useful for improving wayfinding in virtual worlds. Related research will be discussed below. As noted above, Passini defines wayfinding as a three-part process - cognitive mapping, planning, and plan execution; electronic worlds might facilitate wayfinding by supporting each of these processes explicitly. Finally, Passini suggests that wayfinding information be available on a timely basis during decision planning; this suggestion supports research on information scent or residue (trace) (Pirolli and Card, 1995; Furnas, 1997).

2.4.3.2 Virtual Worlds

Several researchers have adapted these physical wayfinding principles to virtual worlds, with mixed results. Two studies by Darken and Sibert (1993, 1996) support this adaptation, but a study by Satalich (1997) does not. An overview of human factors issues in virtual environments in general is provided by Stanney, Mourant, and Kennedy (1998).

Darken and Sibert (1993) investigated the problem of VR users' maintaining knowledge of locational information, that is, current position and orientation. In addressing this problem, the authors propose a basic taxonomy of virtual worlds. Virtual worlds have three primary attributes - size, density and activity - each of which has a continuum of values. As part of their study, Darken and Sibert implemented a tool set of techniques based on physical navigation: flying, spatial audio, bread-crumb (history) markers, coordinate feedback (grid), districting, landmarks, grid navigation (Cartesian or polar), and map views (either forward-up or north-up).

Darken and Sibert (1993) conducted an informal study on the effects of this tool set. The experiment used a simple, sparse, large world randomly configured for each trial. The subjects performed three types of search: exploratory (without target), naïve (exhaustive), and primed (non-exhaustive and informed). Nine subjects each completed a three-part task per trial: explore the space, find the target object, and return to the home position. During the experiment, participant behavior and comments were recorded in written notes by an evaluator. Results generally showed that subjects used environmental cues to partition spaces for exhaustive search (with respect to some arbitrary granularity). Moreover, orientational cues were most effective when stationary or moving predictably, and also globally visible. Users were generally able to combine cues from multiple sensory modalities to assist with target searching. Finally, *tools were found to influence user behavior substantially more than did individual differences*. On the basis of this study, the authors concluded that principles from physical navigational aids (e.g., maps) seem to apply to virtual environments.

Darken and Sibert (1996) later conducted a more involved experiment that applied physical wayfinding and environmental design principles to virtual worlds. Following earlier wayfinding research, the authors propose two types of general design principles – organizational and map. Organizational principles, first, should provide the structure for an observer to mentally organize the environment as a spatial hierarchy for wayfinding: divide the world into parts, preserving a sense of "place", organize the parts under a simple principle, and provide frequent directional cues. Map principles, second, should present spatial information in a "flexible, orientation-independent representation of the environment": show all organizational elements and principles, always show the observer's position, and orient the map forwards-up with respect to the observer. As above, the authors classify wayfinding tasks into three types of search: exploratory, naïve, and primed.

Reflecting these principles, Darken and Sibert studied complex searching in immersive virtual worlds with differing environmental cues. The four experimental treatments were as follows:

1. control - no wayfinding assistance
2. grid - organizational principles only (radial grid)
3. map - map principles only
4. map with grid - map and organizational principles

In the experiment, subjects performed map sketching and thinking aloud. Each treatment was recorded with video and audio tape, as well as virtual position sampling. Ten subjects were tested in five large, sparse, static VEs showing sea and land masses. Each treatment required five exhaustive searches for targets, followed by a primed search for the home position. As in the earlier study, results showed that users' strategies depended heavily on environmental cues. In the control condition, the lack of directional cues and spatial organization led to ineffective searching and frequent disorientation. In the grid treatment, the grid supported searching, but it required work for users to maintain orientation. In the map treatment, the map offered a geocentric perspective to enhance the user's egocentric viewpoint; the map also encouraged the use of landmarks and search optimizations. Other general conclusions are as follows:

- Without adequate directional cues, disorientation hinders wayfinding and spatial learning.
- A large world without explicit structure is difficult to search exhaustively.
- A conceptual coordinate system is often imposed on the world as an implicit divider. Some such structure is required for organized exhaustive searching.
- Path following is natural, and users often treat environmental features as paths.
- A map supports search optimizations, since it supplements survey knowledge.
- Dead reckoning is intuitive and natural.

On the basis of this study, the authors' general conclusion is that *physical environmental design principles can be effectively applied to virtual worlds*. Moreover, human conceptions of virtual spaces may be analogous to conceptions of real spaces. Future research is needed in VEs with different spatial features.

Research by Satalich studied navigation and wayfinding in immersive VR, with the goal of designing tools and environmental cues to enhance navigational awareness. Several questions drove this research. First, is navigational awareness best obtained by environmental self-exploration, or by active or passive guidance? Second, what tools most benefit navigational awareness? Third, how do map study and/or direct navigational experience affect later wayfinding in the same environment, and what is the effect of using a map during exploration? The first part of this question followed the research of Thorndyke and Hayes-Roth (1982). To answer these questions, Satalich conducted a study with 65 subjects. The virtual environment (VE) represented a building 100 feet square, with a ceiling height of 10 feet and 39 rooms. The building contained 500 objects to be used as landmarks. The experiment had a 3x2x2 between-subjects design. The first factor was the type of exploration, which lasted 30 minutes: self-exploration, in which subjects freely explored the building; active guidance, in which subjects followed a pre-determined path, and passive guidance, in which the participant was moved by the system along a pre-determined path. The second factor was the access to a map for five minutes before entering the building. The map showed the building's configuration and the participant's current position with a north-up orientation. The third factor was the access to a map during environmental exploration. A control group didn't explore the building, but only used the map. Experimental tasks included directional pointing (orientation) to familiar but out-of-sight objects; route and Euclidean distance estimations; and sets of naïve and primed wayfinding tasks.

Results showed that people in the VR training condition performed worse than those with map training. Several factors may explain this result. First, subjects with VR training might have equaled subjects with map training, had they been given more time (e.g., days). Second, the VR interface may have distracted novice users, which longer training time would have minimized. Third, the awkward VR hardware interface may have prevented subjects from engaging in natural wayfinding behavior. Further results indicate that maps may have interfered with learning during exploration, perhaps by distracting users from environmental cues. According to Satalich, the most important conclusion is that differences between real and virtual environments affect performance in simple navigation and wayfinding tasks. Future research can clarify these differences, in order to improve the utility of VR for training tasks.

The virtual wayfinding research discussed above reveals striking, but not irreconcilable differences. Satalich acknowledges the existence of several potential confounding factors in her study, which may have obscured wayfinding issues. The studies by Darken and Sibert kept technological and timing conditions constant. Their design may help to explain the effective adaptation of real-world wayfinding and environmental-design principles. Furthermore, although the Satalich VE was smaller, the Darken and Sibert one was topologically simpler. In the simpler environment, users may have found the unfamiliar VR hardware easier to handle. Similarly, the Satalich study included a number of low-level tasks (e.g., distance estimation); the Darken and Sibert study focused instead on high-level wayfinding tasks. This high-level focus may have reduced the impact of hardware considerations. Finally, Satalich's environment displayed route information (paths) under some conditions; Darken and Sibert's environment displayed survey information (grid) under some conditions. As an organizational principle, survey information has greater utility for inferring new routes; this difference may help to explain the experientially-trained Satalich subjects' relatively poor performance, and subjects' relatively successful wayfinding under Darken and Sibert's grid-condition. Hardware and user differences may also have played a role. In summary, under certain conditions, visual fidelity was found to be less helpful than expected. Ultimately, as Satalich suggests, virtual reality is only an analogy or metaphor for the physical world. The relationship between these two domains remains a central issue in navigational research.

Related research was carried out by Elvins, Nadeau, and Kirsh (1997) and Elvins, Nadeau, Schul, and Kirsh (1998) on 3D "thumbnails" for wayfinding in virtual environments. The work explored a multimodal strategy for integrating 3D environmental cues with UI components for navigation and manipulation. Traditionally, such UIs feature text and/or 2D images. The authors designed and developed a prototype using miniature 3D representations of sections of a virtual environment, in order to support the user interface. These representations were termed “worldlets .” A pilot study and later experiment showed that worldlets improved users’ landmark knowledge and expedited wayfinding in virtual environments. The long-term utility of this prototype widget is unclear, as is its potential integration into 3D toolkits for users. Yet the research links two often-distinct software domains – the 3D model and the 2D controller – and explores innovations in design and evaluation.

Another study of multimodality in virtual environments was conducted by Czerwinski et al. (1999). Using a desktop 3D application called the Data Mountain, participants had previously stored 100 Web page bookmarks - thumbnails and titles - in locations of their own choosing. Four months later, the participants were tested for ability to retrieve those bookmarks, with or without the thumbnail information visible. Mouse-over title text was always available. Results showed that subjects experienced almost no degradation in retrieval ability with thumbnail information visible, and similar performance (after two out of five blocks of trials) with thumbnail information invisible. The study shows the power of voluntary spatial placement (and grouping) for later cueing of related semantic information. Future work might well examine scalability of the phenomenon, as well as the effects of mixing voluntary and/or involuntary information placement in virtual space.

Clarifying some issues of navigational tools and human psychology, Wickens (1992) reviews research in spatial perception and cognition and the display of spatial information. Key topics include the acquisition of navigational information and spatial representations of non-spatial systems. As discussed previously, people acquire navigational knowledge in three stages: landmark, route, and survey (McKnight et al., 1991; Thorndyke and Hayes-Roth, 1982). This acquisition process has implications for training and other tasks. Route knowledge is ego-centered, with the advantages of automaticity and compatibility with the frame of reference. Survey knowledge, by contrast, is world-centered, with the advantages of flexibility and superior decision-making. Route lists are more verbal, egocentric, and compatible with the frame of reference. Maps, by contrast, offer superior support for navigational decisions. Rotating maps feature more congruence with the frame of reference, but fixed maps offer more consistency and portability. (A potential compromise is a fixed map with a moving point-of-view indicator.) For spatial representations of non-spatial systems, Wickens refers to four guidelines for the *visual momentum* principle:

1. use consistent representations
2. employ graceful transitions
3. highlight anchors (the display’s invariants)
4. display continuous world maps

Vinson (1999) extrapolated design guidelines for virtual environments from the research above and related work. The guidelines focused on Lynch's (1960) design elements to support wayfinding. Vinson reiterates the importance of these elements, while translating them to virtual environments. Such elements should be visible and distinct from neighbors, as well as featuring elements to distinguish them from data objects. Placement of design elements on paths and their junctions is useful, Vinson proposes, as is the use of a grid for placement and alignment. This research systematizes a body of work not widely applied to VEs. Implementation and evaluation have not been reported to test the guidelines empirically.

Wickens' comments apply to the research of Darken and Sibert (1996) and Satalich (1997), as well as to future work in this area. According to the principle of visual momentum, the VEs of Darken and Sibert, and Satalich, were good spatial representations of non-spatial systems. These VEs followed the first three guidelines of the principle, and the fourth guideline became an experimental condition. Given the difficulty of orientation during the wayfinding tasks of object search, the maps of Darken and Sibert were appropriately designed as rotating maps. The fixed map of Satalich may have been less appropriate for these tasks. Wickens' comments further suggest that the guided paths of Satalich were practical for route-following, but not useful for building survey knowledge. In general, Wickens' work suggests that designers of electronic environments and tools should consider the tradeoffs between route and/or survey knowledge with careful consideration of users' key tasks.

2.4.4 Hypermedia Environments

In hypermedia as well as in spatial worlds, there are issues of information exploration. The discrete nature of hypermedia, however, raises some new issues. This section will first consider these issues from a theoretical perspective, and then report on several characterizations of user behavior in hypermedia.

Citing "experiential, anecdotal, and empirical" evidence of disorientation, McKnight et al.(1991, 1993) conducted research with text and hypertext documents. Disorientation is described as users not knowing where to go next, or not knowing their location in the document structure. The authors accept the predominant metaphor of navigation. For electronic information spaces, they propose a four-level psychological model: schemata, landmarks, routes, and surveys. People develop schemata or models of the physical environment through experience, for the purposes of orientation and navigation (Downs and Stea, 1973). During learning, people successively instantiate schemata with landmarks, routes, and surveys (Thorndyke and Hayes-Roth, 1982). For print information, documents generally have well-known schemata "that facilitate comprehension of material by allowing readers to predict the likely ordering and grouping of constituent elements of a body of text" (McKnight et al., 1991). Although navigational principles might apply, the analogy with physical navigation breaks down beyond landmark knowledge.

Route knowledge is generally not needed in a text with random access, and survey knowledge generally covers a text's contents (not its form). Usually newer and more diverse in organization, electronic documents lack standards and transparent structures. Unlike books and newspapers, hypertext documents reveal little information at a glance (e.g., size, quality, age, past usage). Users' schemata of hypermedia are probably more abstract and less substantive than those of paper documents. Hypertext users rely heavily on landmarks, but the acquisition of further navigational knowledge is not well understood. To facilitate user navigation in electronic documents, the authors recommend graphical browsers, maps, and structural cues.

Given the restricted two-dimensionality of hypertext structures, the absence of strong user survey knowledge is not surprising. Although not mentioned by the authors, anecdotal evidence suggests strong user route knowledge, which is supported by the topology of hypertext. Research is needed to determine the role of survey knowledge in hypermedia, as well as the design of appropriate maps or other tools. Hypermedia often has a semantic, rather than physical structure; the implications of this issue will be discussed below. Meanwhile, in the absence of standards and familiar structures, the adaptation and presentation of useful schemata for hypertext seems a fruitful area for research. (A general model that might support hypermedia database systems is proposed by Tompa, 1989.)

A review of research on spatial metaphors and disorientation in hypertext browsing is provided by Kim and Hirtle (1995). The authors identify three classes of disorientation problem:

- *navigational* – inadequate knowledge of information organization and tools
- *"art museum"* - high-quantity, low-quality information assimilation
- *embedded digression* - confused task switching

Similarly, a hypertext user's cognitive load includes three types of task:

- *navigational* - wayfinding
- *informational* - database analysis and summary
- *management* - of navigational and informational tasks

Spatial metaphors will generally assist with wayfinding tasks, but these metaphors won't facilitate informational tasks and task management. Thus two complementary types of tool are needed to reduce disorientation – for general wayfinding or for specific tasks.

User navigation in complex database structures is characterized by Canter, Rivers, and Storrs (1995). In this work, the authors apply psychological concepts of navigation and path algebras. This research considers hypermedia links, UI control options, and task constraints. The authors developed six indices to characterize user navigational behavior:

1. *pathiness* (PQ) - a path is a route that crosses no node twice
2. *ringiness* (RQ) - a ring is a route which returns to the starting node
3. *loopiness* (LQ) - a loop (circuit) is a ring containing no sub-rings
4. *spikiness* (SQ) - a spike is a route that returns to its origin by retracing the visited nodes
5. *NV/NT* - the ratio of nodes visited (NV) to nodes total (NT)
6. *NV/NS* - the ratio of different nodes visited (NV) to the total number of visits (NS).

Experiments were performed with a network-structured database system. On the basis of observations, the authors describe five information-seeking strategies to characterize user navigation topologically:

- *Searching* features ever-increasing spikes with some loops, as users seek a specific target.
- *Browsing* features many long loops and some large rings, as users wander until their interests are caught.
- *Scanning* features a mix of deep spikes and short loops, as users cover a wide area shallowly.
- *Exploring* features numerous paths, as users survey the extent and nature of the data.
- *Wandering* features many medium-sized rings, as users navigate in an unstructured way.

In conclusion, Canter et al. suggest the value of identifying and clarifying data landmarks (as do Valdez et al.1988; Ingram and Benford, 1995). The authors' research suggests the value of characterizing users' navigational strategies in hypermedia such as the WWW and in virtual environments; this characterization could be used to design Web sites or virtual worlds to support the strategies associated with planned user tasks. The relative influence of user differences, task, and hypermedia or virtual topology is worth investigating. Finally, the navigational strategies of Canter et al. recall the optimal foraging models of Pirolli and Card (1995); further research might investigate this potential connection.

Focusing on differences between browsing and searching, Campagnoni and Ehrlich (1989) conducted a study of users' strategies and skills for IR in a hypertext help system. The study had a psychological focus, rather than the structural emphasis of Canter et al. Two aspects of the study by Campagnoni and Ehrlich will be discussed here – information-seeking strategies and spatial visualization ability. Developed for Sun workstations, the hypertext help system features a three-level hierarchy of handbooks, as well as a master index. The system employs a page-oriented hypertext database, with a GUI front end. Twelve subjects acted as system administrators to answer six user questions. Data were collected by software monitoring and videotape recording. After the session, each participant was given a standard test of visualization ability. The study defined two strategies. The browsing strategy involved scanning tables of contents and paging through relevant topics; the analytical strategy (searching) involved using indices to look up terms and then following links to topics and pages. Despite questions designed to elicit both strategies, most users preferred browsing. In the help system's shallow hierarchy, browsing incurred no time penalty to discourage its use. Inexperienced users sometimes browsed the right sides of index pages; such users could not formulate effective queries for searching. Predictably, searching was performed most often by expert users, and as a last resort by browsing users. The study found a significant negative correlation between visualization ability and total question-answering time. Presumably, users with good visualization skills construct superior mental models of the system's information architecture, which support effective navigation; these models also help to prevent disorientation in the absence of visual cues for information organization. The authors note the *tradeoff between the low cognitive load of browsing and the power of searching. They also note that systems can bias users towards an information-seeking strategy* by search result list length, link and index quality, UI design, and database size and type.

The Campagnoni and Ehrlich study showed a strong user preference for browsing, rather than searching. This preference underlines the importance of good navigational support in hypermedia systems. Unfortunately, users with relatively weak visualization skills perform badly on navigational tasks. Developing non-spatial navigational tools is important for this user group. Two research avenues suggest themselves. First, route knowledge could be emphasized over survey knowledge, since Wickens (1992) noted that route knowledge is amenable to verbal description. Second, semantic models could be proposed to users, rather than physical ones; this distinction is discussed below.

In a large study, Catledge and Pitkow (1995) conducted an experiment to characterize browsing strategies in the WWW. For the study, the authors captured UI events in a Web browser at a university laboratory for three weeks. The study included more than 100 users, 1200 Web sites, and 43,000 UI events. Results showed that hypermedia links were preferred for node traversals (52%), although the browser's "Back" button accounted for almost as many traversals(41%). The average number of documents requested within a site was 12.64. For external sites, an inverse relationship was found between access frequency and average navigational path length per visit. The authors characterized users as follows:

1. *serendipitous browsers* - avoid repeating long paths and browse shallowly
2. *general-purpose browsers* - have an average (0.25%) chance of repeating a complex navigation sequence
3. *searchers* - seldom repeat short navigational sequences, but follow long sequences often.

In terms of the research by Canter et al. (1995), serendipitous browsers are wanderers; general-purpose browsers are scanners, browsers, and explorers; and searchers are searchers. Overall, users tended to remain in a small area within a site; their navigational paths resembled a hub with spokes, on account of frequent backtracking. Users rarely traversed more than two layers in a hypertext structure before returning to a home point. Extracting design guidelines, the authors recommend that important information be accessible within two to three jumps from a user's home page. Also, frequent document indices support the observed hub-and-spoke navigational pattern, as well as reducing user disorientation. Finally, document designers should expect different classes of user (as described above) and perhaps create distinct documents or views for each class. In particular, there is a trade-off between the "volatility" enjoyed by a browser and the efficiency required by a searcher.

In a related study, Tauscher and Greenberg (1997) studied users' revisitation patterns in Web navigation. The researchers sought to gather empirical data for the design of history mechanisms in Web browsers. For six weeks, the browsing data of 23 users was gathered. Results showed that 58% of document accesses were revisits, and that users continually expand their document access set. Also, users often returned to pages recently visited; they accessed relatively few pages often; they tended to browse in small clusters of related pages (a working set), and they repeated only short node paths. With regard to browsing mechanisms, the stack-based history method of commercial browsers was found to be inferior to showing the last few recently visited nodes, with duplicates removed.

Both the Catledge and Pitkow (1995) and the Tauscher and Greenberg (1997) studies showed strong user homing behavior. Both studies also showed strong locality of reference away from the home position. Moreover, the absence of long repeated paths (shown by Tauscher and Greenberg) suggests a frequent need for users to infer new routes. For these three reasons, hypermedia system support for global or dynamic, local versions of survey information might improve the quality of user exploration. Related research has been conducted in this area (e.g., Andrews, 1995; Mukherjea et al., 1995).

2.4.5 Spatial Ability in Electronic Environments

In considering electronic environments of any type, it is important to consider individual differences in user abilities. There are two general views on this issue (Chen, 2000): one sees differences as reducible through education and training, while the other sees differences as hard to change, but manageable through specially-designed tools. In any case, in HCI, a core set of basic cognitive abilities has been identified as influencing task performance in predictable ways. Psychological measures of such abilities not only establish research context, but also help to generalize research findings. Abilities relevant to navigation in information structures, in particular, include spatial ability, associative memory, and visual memory.

Spatial ability in information exploration has received significant research attention. Chen (2000) cites several studies confirming the importance of this ability; one such study was that by Campagnoni and Ehrlich (1989) mentioned previously in this chapter. In these cited studies, spatial ability has generally been found to correlate positively with task performance in information exploration tasks. Two spatial-ability tests used in such research are (1) the kit of factor-referenced cognitive tests by the ETS in Princeton and (2) the Minnesota Paper Form Board Test by the Psychological Corporation (Likert and Quasha, 1995). The Minnesota test has been widely used for occupational studies and "results indicate that it is one of the most valid available instruments for measuring the ability to visualize and manipulate objects in space." (Anastasi, 1988). The test has also been employed in research on information exploration, e.g., an evaluation of students in a virtual teaching environment at the University of Edinburgh (Cronin, 1998).

Chen and Czerwinski (1997) and Chen (2000) conducted empirical studies of the relationship between individual differences and information exploration in a virtual world organized on the basis of visualized semantic relationships. As a general goal, the studies sought to assess the usability of a prototype VE, and to examine methodological issues for development. In so doing, the studies examined the search strategies and general preferences of users. Three cognitive abilities mentioned above – spatial ability, as well as associative memory and visual memory – were central to the studies. For assessing individual ability, the researchers used the ETS kit mentioned above. This research was partly informed by work of Darken and Sibert (1996) and Lynch (1960), to be discussed later in this chapter.

The virtual environment for the studies was an early version of the StarWalker software (Chen and Czerwinski, 1997; Chen, Thomas, Cole, & Chennawasin, 1999). The underlying information domain was a collection of 169 papers from the ACM CHI conferences of 1995, 1996, and 1997, visualized as a network. As a first step in visualization, papers' pair-wise content similarity was computed by latent semantic indexing (LSI), an information-retrieval technique. The strongest such similarities were extracted as a matrix using a type of network scaling that provides particularly accurate details of local structure. Finally, the resulting semantic space was visualized in VRML (Virtual Reality Modeling Language), a standard language for desktop virtual reality. Users could move freely through the virtual environment, with view zooming to mediate between overview and detail. The document collection was thus presented in semantically-related clusters in the virtual environment. Each paper was represented by a virtual sphere, whose color indicated the year of publication; author initials annotated each sphere. In the user interface, the left panel contained this virtual environment; the right panel held a textual frame that displayed paper abstract information when a user clicked the mouse on a virtual sphere. A document search resulted in cylindrical spikes being displayed above the 20 most relevant document spheres (after the first study); these spikes served as visual landmarks.

The first study (Chen and Czerwinski, 1997) involved 11 participants. A pre-test was administered for spatial ability, followed a short demonstration of the VRML user interface. Participants were then given 10 minutes to find as many papers as possible in the virtual environment that were related to a particular topic; participants were then asked to sketch the virtual environment from memory. Participants were next given 10 minutes to find five papers on a given topic; afterwards, they were asked to categorize and abstract items in the semantic space by naming clusters of papers. A post-test questionnaire asked about software usability and user satisfaction.

Results showed that spatial ability was positively correlated with recall (number of relevant abstracts found, as a share of the total number of relevant abstracts available), but negatively correlated with precision (number of relevant abstracts found, as a share of the total number of abstracts found). Spatial ability was also positively correlated with the number of abstracts judged relevant by each participant. These results suggest a connection between spatial ability and ability to utilize the structure of visualization. Participants' world sketches varied substantially, whose accuracy apparently reflected individual differences in spatial ability. Results also showed that subjects retrieved more abstracts that were located near to structural joints in the visualization, relative to abstracts located farther from such joints. Moreover, subjects tended to navigate slowly outwards from the center in their search patterns. In general, participants liked the virtual environment, but identified usability problems such as unfamiliarity with VRML viewers and clustering models for visualization. In addition, the researchers noted the need for a larger number of explicit navigational cues.

The second and third studies (Chen, 2000) followed the general methodology of the first one, while focusing on associative memory and visual memory, and associative memory and spatial ability, respectively. In the second study, a positive correlation was found between associative memory and overall recall. In the third study, precision for users of a spatial UI was negatively correlated with associative memory, while precision for users of a textual UI was negatively correlated with spatial ability. In general, associative memory seems a good predictor (positive or negative) of information foraging performance. Moreover, participants with good associative memory tended to prefer a spatial UI to a textual one. From a design perspective, enhancing visual cues in the virtual environment might help to compensate for individual differences in this regard, by helping users to identify and find virtual structures. The three studies did not generally find a significant, general impact of spatial ability on user performance, though particular results (as mentioned above) suggest that further research in this area would be worthwhile (e.g., Waller 1999, discussed below).

Later research extended this virtual environment under the name of StarWalker (Chen et al., 1999). The new environment permitted social navigation, in the style of online virtual communities, complete with avatars (virtual personae) and textual chat. The research goal was to unify spatial models, semantic structures, and social navigation. In particular, the researchers were curious about the effect of virtual world organization on users' communication patterns and styles. Preliminary results showed that the environment supported focused and articulate communication on topics relevant to the environment's data.

As suggested above, further research would be beneficial in the area of spatial ability and virtual environments. Waller (1999) performed an assessment of individual differences in spatial knowledge of real and virtual environments. Results showed that the variation in spatial knowledge of virtual environments was even more substantial than the already-large variation in spatial knowledge of real environments. For this reason, Waller suggests that VE research in this area consider individual differences explicitly, as they may be more significant in accounting for experimental variance than design differences. Such consideration can be given either by testing for individual abilities in conjunction with evaluation of 3D environments, or by controlling for individual differences statistically.

In Waller's (1999) study, the largest single contributor to experimental variance was proficiency with the user interface, followed by spatial ability and then gender. The significant impact of UI competence suggests that studies in virtual environments should allow sufficient time and resources for user training, in order to minimize the effect of this competence. Spatial ability had relatively less impact on performance, but this ability still predicted VE spatial knowledge acquisition, though not real-world spatial knowledge. It appears that users of desktop virtual environments didn't need to know their spatial locations (as required when navigating in the real world), but instead required an ability to understand 2D spatial information and perspective geometry. Nevertheless, a maze-based training task that Waller conducted in a VE proved somewhat predictive of later real-world navigation based on this maze. Finally, with regard to gender, this factor was shown to be predictive of performance on a task basis, mainly through correlation with interface proficiency and spatial ability. Overall, the three individual differences that contributed to understanding task performance in Waller's study – UI proficiency, spatial ability, and gender – accounted for at most half of the experimental variance. Further research could extend this understanding of the role of individual differences in virtual environments.

The research discussed in this section is important in that it unites themes from throughout this literature review – information structuring, exploration, and visualization. The work combines a variety of techniques from the applied and social sciences to investigate both personal and social navigation in novel ways. Given the novel nature of the research domain, however, the empirical research is unavoidably inconclusive. Future work would be useful with regard to spatial ability, as mentioned previously, as well as generalization of previous work to other information tasks (e.g., without querying) and structures (e.g., hierarchical data).

Certain refinements in the experimental methodology (Chen and Czerwinski, 1997; Chen, 2000) might also be useful. The VR training time could be extended, perhaps to 30 minutes, so that participant response would more fully reflect the data model and information tasks (Waller, 1999). Second, a single-pane display could be used. (The prototype's multi-pane display might have introduced a confounding factor, in requiring the user to switch attention between spatial and textual styles of interface.) Third, it would be useful to introduce some sort of click-and-jump navigation to ease the burden of mouse and keyboard controls for complex 3D navigation. Finally, it is important to investigate whether more 3D input/output hardware would be more effective in presenting 3D information models, rather than on a 2D desktop computer display.

2.4.6 Discussion

To sum up this section, then, research on physical and information exploration has considered a variety of issues with a range of approaches. A few themes and generalizations stand out among the variety:

- Navigational expertise consists of an expanded knowledge of domain structure and environmental perceptual cues. Well-designed environments can support this expertise.
- Hierarchy is fundamental to psychological and information structures.
- The distinctions between landmark, route, and survey knowledge are fundamental. They appear in the physical world, as well as in both hypermedia and virtual worlds. Accordingly, environment and tool design should consider these distinctions.
- Wayfinding and design principles from the real world seem to apply to electronic worlds, with certain qualifications and adaptations.
- A successful study of navigation should use several signatures (sets of cognitive encoding and decoding operations) and data collection mechanisms.
- Environment, tools, and individual differences (e.g., in spatial and memory abilities) affect user strategies and performance, to varying degrees.

2.5 Information Visualization

The information to be assimilated falls broadly into two categories, identified by Benedikt (1991) as navigational ("information that serves to organize us and the world in spatiotemporal terms") and destinational ("information judged to be of intrinsic value"). The bulk of this review so far has focused on research into navigational issues, as users struggle with potential disorientation in complex information worlds. Issues with destinational information have been subjected to substantial research as well, as theories have been proposed and techniques developed to display facts and relationships in comprehensible ways using information systems. The process of presenting online destinational information in this way has generally been termed information visualization. As globally networked information spaces develop, information visualization research has begun to consider navigational issues as well. (Besides, as mentioned above, multi-purpose elements can be designed to serve as both navigational and destinational data.) A general definition of information visualization is the following: "The use of computer-supported, interactive, visual representations of abstract data to amplify cognition." (Card, Mackinlay, & Shneiderman, 1999). "The fundamental strategy of visualization is to convert data to a visual form that exploits human skills in perception and interactive manipulation." (Card et al., 1999).

This section reviews key research in the field of information visualization: fisheye views, semantic zooming, dynamic queries, XEROX PARC's 3D tools, information landscapes and spaces, hypermedia structure viewers, and dynamic systems. Although previous sections have discussed both physical and electronic worlds, this section will naturally focus on electronic worlds. Overviews of this area are provided by Chen (1999) and Card et al. (1999).

In reviewing research in information visualization, a series of theories and designs will be considered. Several concepts and issues, some discussed earlier in this review, will be used to consider these visualization approaches:

- design of a hypermedia or spatial world
- type of information structure
- technique for integrating physical and semantic structures
- utility for wayfinding - mental mapping, route planning, or plan execution
- data type - destinational or navigational
- utility for electronic environments or maps
- representation strategy for values or structure - raw or derived (Tweedie, 1997)
- presentation attributes, such as static or dynamic organization, simple or compound display, two or three dimensions, and textual or graphical presentation

2.5.1 Fisheye Views

In information systems, a general problem is the lack of display space to show all available information. Furnas (1986) proposed generalized fisheye views as a solution to this problem. This technique mimics the perceptual structure of the human eye: it displays local detail and global context simultaneously. While users generally require content detail to interact with information structures, users often become lost and lack interpretative contextual information. Fisheye views solve this problem by trading off *a pri-*

ori importance against visual distance (or prominence). Formally defined, a fisheye view is created using a Degree-Of-Interest function composed of *a priori* importance and distance measures. These views can be implemented for a wide range of information structures, particularly trees, but also lists, DAGs, graphs, and spaces. In a usability experiment, generalized fisheye views significantly outperformed traditional flat hierarchical views.

Sarkar and Brown (1992) proposed an extension of fisheye views for 2D graphs and other structures. They enhanced the fisheye technique with layout considerations, including object positions, and detail size and level. These layout attributes were computed on the basis of an object's distance from the user's current focus, as well as the object's pre-defined importance in the global information structure. Sarkar and Brown assert that their enhancements make fisheye views more expressive and natural. They developed a prototype system, which can maintain real-time response for graphs containing 100 vertices and 100 edges. The authors noted potential extensions to their work for multiple viewing foci and slave viewers for node content.

Keshkin and Vogelmann (1997) developed an algorithm for visualizing hierarchical graphs with a 3D landscape metaphor. Related to research on fisheye views (Sarkar and Brown, 1992) and the Information Visualizer (Card et al., 1991; Robertson et al., 1991) described in a later section of this chapter, the algorithm lays out hierarchical information in a plane. Child nodes in the hierarchy are placed recursively within the area allocated for each parent node. Distance and object scale allow the user to zoom hyperbolically into or out of the data. The algorithm is simple, yet effective for structuring hierarchical information landscapes.

Fisheye views can be developed for hypermedia and spatial worlds, either two- or three-dimensional. They do not alter the values or structure of presented information, but they use both physical and semantic information to compute visual distance or prominence during fisheye distortion. The technique facilitates mental mapping by maintaining continual context for detail views, but it can distort the distance judgments needed for developing survey knowledge and cost-effective route plans in spatial worlds. The technique displays data that is both navigational and destinational. For navigation, the technique may be ideal for maps, and useful for building environments.

2.5.2 Semantic Zooming

Also seeking a balance of overview and detail, another graphical technique is known as semantic zooming or multi-scale interfaces. A physical zoom, on the one hand, changes the size and visible detail of objects. A semantic zoom, on the other hand, changes the type and meaning of information displayed by the object. (These techniques can be combined.) Semantic zooming avoids the physical distortions of fisheye views, by using a semantic transition between detailed and general views of information. Semantic zooming is exemplified by physical maps with different features and organization at different levels of scale. Furnas and Bederson (1995) have proposed a general analytic framework for multi-scale interfaces. Their “space-scale diagrams” represent a spatial world and its possible magnifications, which allows the analysis and visualization of scale-related UI issues.

The Pad project developed a prototype interface based on semantic zooming (Perlin and Fox, 1993). Pad features a multi-user, infinitely wide 2D workspace. Designed to take advantage of human spatial and geographical skills in navigation, Pad seeks to use visual mental mapping to organize large information spaces. For navigation, Pad features mobile “portals” or teleporters. These portals support viewing and transportation links to non-local parts of the workspace. These portals are so-called “magic features,” and they function as hypertext links (Dieberger, 1996). As local and remote views change, objects display semantic detail in accordance with distance from the user. Sample applications include hypertext editing and browsing, which treat links as detail items into which to zoom.

Bederson and Hollan (1994) designed a successor to Pad, named Pad++. Unlike Pad, Pad++ is conceived as an application substrate. The project targets visualization and browsing in information-intensive domains, such as hypertext, computer file systems, and historical timelines. The authors proposed zooming as part of an “interface physics,” which uses physical models for the visible behavior of objects, as suggested by Benedikt (1991). Interface physics is presented as an alternative to higher-level UI metaphors, which may not scale well and sometimes raise false expectations. Navigation in Pad++ occurs in implied parabolic jumps, as the interface zooms out of a location to show context, and then zooms in to a new location to show detail. Although potentially slow and intrusive, this technique maintains navigational context and may facilitate the development of survey knowledge. A sample application showed geographical maps at several levels of scale in an intuitive fashion.

Semantic zooming can be developed for hypermedia and spatial worlds, with a variety of information structures. It does not change the values or structure of an electronic world, but it uses semantic information to change the physical representation of objects according to viewing scale. Semantic zooming is likely to facilitate mental mapping and route planning, by automating the hierarchical representation that underlies these processes (Chase, 1983; Passini, 1984; Stevens and Coupe, 1973). The technique displays data that is both navigational and destinational. For navigation, the technique is suited to both environments and maps, particularly the latter (Bederson and Hollan, 1994). The technique can be used in two- or three-dimensional worlds.

2.5.3 Cone Tree, Hyperbolic Browser, and Butterfly

Adopting a related approach to balancing overview and detail are three visualization designs (and associated theoretical principles) from XEROX PARC. This project's complex, non-immersive, spatial worlds; explicit tradeoffs between perceptual and cognitive tasks; and innovative visualization designs have implications for navigation in electronic worlds. This section will first discuss general principles and goals, then describe specific designs, and conclude with some observations on navigational issues.

Often facing information overload, knowledge workers are limited in the amount of information that they can manipulate; they often have difficulty detecting patterns in, and deriving meaning from, information. To address this problem, researchers at PARC have developed an information workspace known as the Information Visualizer (IV) (Card et al., 1991; Robertson et al., 1991). The IV provides a framework for several visualizations that incorporate computer graphics, systems architecture, and cognitive psychology. A key design principle is "to shift some of the user's cognitive load to the human perceptual system," thereby enabling faster task processing (Robertson et al., 1991).

The project's UI design goals are derived from six observations about information processing systems:

- the benefit of hierarchical system organization
- the high cost variability of information storage, for both finding and assimilating information
- the user's tendency to locality of reference in information processing
- the user's tendency to reference clustering in task performance
- the benefit of maximizing the ratio of information to cost
- the benefit of bottom-up aggregation and abstraction in an information processing system

A "focus+context" technique (abstracted fisheye view), the Cone Tree assists with managing and accessing large hierarchical information spaces (Robertson et al., 1991). Cone trees are hierarchies laid out in 3D, as shown in Figures 2.1(a) and 2.1(b). Each set of child nodes is displayed as a circle, which is parallel to other such circles. When a node is selected with a mouse, the tree rotates to the front; and the system highlights the selected node and each node in its parental path. The cone tree uses a vertical layout; a cam tree uses a similar horizontal layout. For hiding selected parts of a complex hierarchy, the cone tree supports the operations of pruning and growing. Dynamic structural modifications are possible with a mouse. IR is supported for visible nodes. Sample applications include browsers for files and organizational structures. Potential applications include software and document management, and WWW browsers.

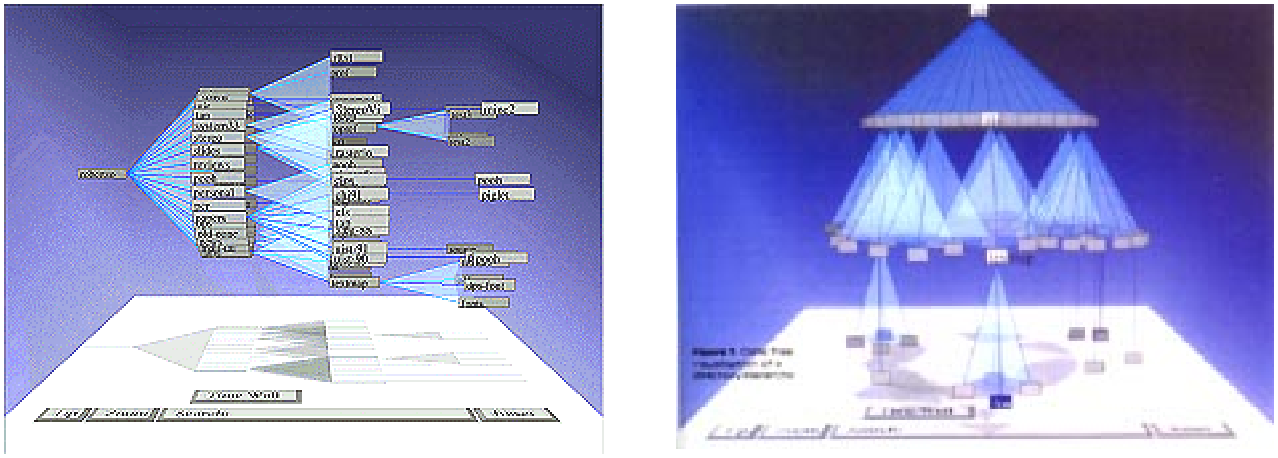

Figures 2.1(a) and 2.1(b). Horizontal and vertical cone trees (Robertson et al., 1991).

Cone trees were used to prototype and evaluate a new user interface that integrated searching and browsing of large category hierarchies and associated text collections (Hearst and Karadi, 1997). The prototype was known as “Cat-a-Cone.” Cat-a-cone supported the simultaneous display of multiple categories in context, as well as point-and-click Boolean queries using the visual representations. Sample data for the research project included a collection of medical texts.

Another visualization tool, the hyperbolic browser, is a radical fisheye visualizer for large hierarchical structures (Lamping et al., 1995). Mathematically, the technique lays out a hierarchy on a uniform hyperbolic plane, and then maps this plane onto a 2D circular display region (as shown in Figure 2.2). Hyperbolic visualizations have two important properties:

- Components diminish in size as they move outwards.
- Moving outwards from the center, there is an exponential growth in the number of components.

So, visual context always includes several generations of nodes, which facilitates user orientation. The hyperbolic browser can handle hierarchies much larger than those of conventional hierarchical browsers. This geometry requires two potential corrections to support intuitive movement: either the root node or the focus node should maintain a canonical orientation with respect to the screen. During usability evaluation with four subjects, no performance differences were observed relative to a traditional hierarchical interface. Users expressed a strong preference for the new interface, both for searching and for learning overall structure.

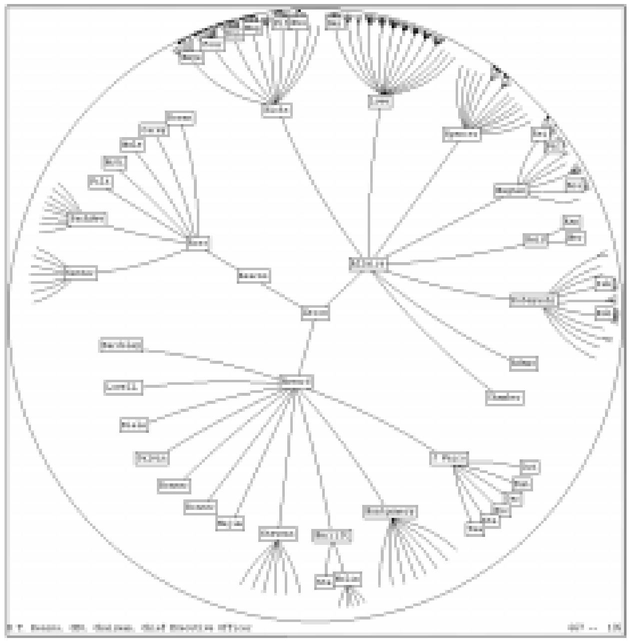

Figure 2.2. Hyperbolic browser (Lamping et al., 1995).

As shown in Figure 2.3, Butterfly is a specialized, compound IV application for accessing DIALOG's Science Citation databases on the Internet (Mackinlay, Rao, & Card, 1995). Butterfly addresses the problem of a fast UI but multiple, slow data sources. Butterfly integrates search, browsing, and access management with four techniques:

1. Visualization of references and citers supports user understanding of retrieved information, and it supports browsing of search results
2. Automatic "link-generating" queries assemble bibliographic records into citation graphs.
3. Asynchronous query processes explore citation graphs on behalf of the user.
4. Process controllers let the user manage query processes.

The animation loop and asynchronous query processes give Butterfly an organic feel. User studies show improved speed and reduced training time, relative to normal DIALOG usage, with some loss of query power. The authors extend this design to a general information access approach called “organic user interfaces for information access.” In this approach, *a virtual landscape is grown under user control as information is accessed automatically*. These proposed UIs have four key components: information landscapes, growth sites, growth agents, and growth controllers.

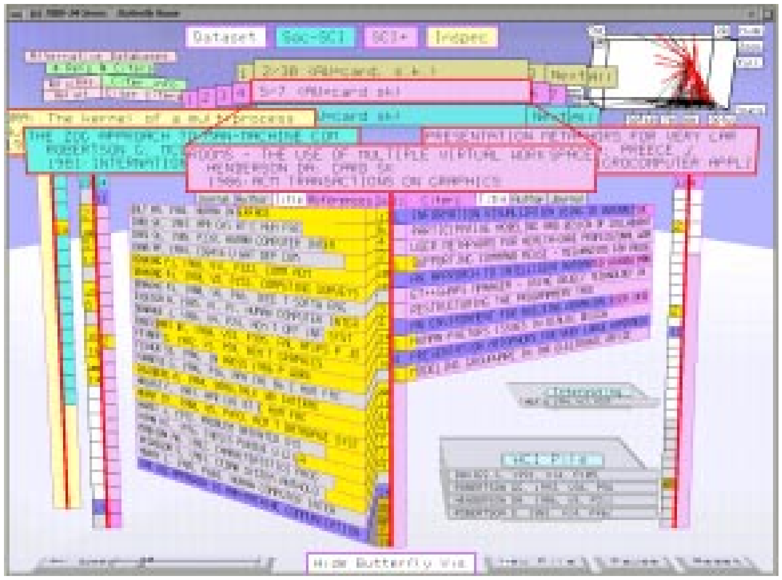

Figure 2.3. The Butterfly application (Mackinlay et al., 1995).

All of the IV tools use non-immersive VR. Each tool is specialized for a particular information structure. While the Cone Tree and the hyperbolic browser use existing values and structure, the Butterfly and the Web tools generate new values and structure. In general, physical and semantic structures are placed in correspondence, as recommended by McKnight et al. (1991). Most of the visualizers are amenable to mental mapping. The IV's general support for pattern recognition and dynamic queries could enhance mental mapping with attributive information. However, the displays of the hyperbolic browser and the Butterfly are probably too dynamic for good mental mapping. Used as tools, several IV features could assist route planning: the Butterfly for cost-benefit analyses, and asynchronous agents for timely information-gathering. Several tools could also facilitate plan execution for following sequences of choices with attributive information. The IV's designs would benefit from behavioral testing for navigational task effectiveness, as well as potential portability to immersive VR. Overall, the IV's combination of low-level physical metaphors (location, motion, perspective, etc.) with high-level domain metaphors (tree, wall, butterfly, book, office) seems well suited to support both locomotion and wayfinding.

2.5.4 Visualization for Information Retrieval

The visualization approaches in the preceding sections have predominantly treated structured data; the approaches in this section will treat predominantly unstructured data. One such approach to the latter involves information landscapes for information retrieval (IR). Despite their power for complex searching in large document corpora, traditional IR systems have several well-known problems from an HCI perspective. These problems include query interfaces requiring difficult languages and without direct manipulation; complex text-only output in the form of document descriptions or lists; concealment of useful modeled information; weak support for viewing object relationships; and generally heavy cognitive demands. Several researchers have addressed these problems with spatial navigational interfaces, as discussed below.

Chalmers investigated user exploration using a map or landscape metaphor. His work derives in part from wayfinding research by Lynch (1960) and Passini (1984). The prototype Bead system had several design goals, including a legible information space (Lynch, 1960), reliance on sensori-motor skills (Card et al., 1991), and good information design. The resulting system builds and displays a spatial model, which shows both detail and overview information. The construction uses patterns of document similarity. Visual proximity thus represents an abstract information dimension, and the spatial metaphor supports visualization of overall relationships. Also, visualized proximity often reveals useful connections. Bead's 2.1 D landscape supports arbitrary points of view. The structure and appearance of the modeled data provide legibility features such as landmarks, districts, and edges (Lynch, 1960). Informal usability testing showed that users preferred keyword-based layouts to random or abstract ones. Keyword searches use color to indicate relevant documents and clusters, as a sort of dynamic query. Users can click on an object to reveal detailed document information. Although meta-information is revealed by landscape's construction, most display dimensions are reserved for information content.

In later research, Chalmers, Ingram, & Pfranger (1996) explored the addition of imageability (legibility) features to the Bead system. With relatively minimal detail, Bead could be made more engaging and legible. The research goal was to improve exploration, navigability, and memorability, without adding detail that would occlude important information. Several features were added to the system. Static features, to improve user orientation, included colored regions, paths, and clusters. Dynamic features included the following:

- sampling to reduce displayed information; this sampling used view boundaries, viewpoint proximity, word frequency, and search histories
- topic labels for regions; these labels used view frequency and historical usage popularity;
- usage disks, which reveal search history

The new system utilizes the history and context of users' activities. This approach simplifies data analysis, and it enriches information visualizations. The Bead system lacks avatars, but it visually represents the activities of all current users. (Related earlier work is reported in Ingram and Benford, 1995.)

LyberWorld, another spatial visualization of a document corpus, supports IR through navigation in complex, abstract information spaces (Hemmje, Kunkel, & Willett, 1994). In the system, user tasks include query construction, orientation in content space, relevance feedback, and orientation in retrieval context. During a spatial search, the user faces two types of task. The first is navigational. The second is informational, which includes evaluating exploration completeness, judging the relevance of retrieved items, and inspecting these items. LyberWorld features three loosely integrated techniques to support these tasks:

- *Navigation cones* are cone trees that contain alternating layers of documents and terms. Users search by manipulating these cones, which provide detail, context and history
- *RelevanceSpheres* use a 3D layout to provide relevance feedback and clustering of retrieved documents. The user can adjust the weights and positions of documents and terms to clarify relevance.
- The *InformationRoom* displays the text of a retrieved document on a virtual wall, for close inspection at the end of the search process.

These three IR visualizations derive spatial structures from document corpus values. In LyberWorld, physical and semantic structures were placed in correspondence. In Bead, the landscape showed physical structure corresponding to global semantic structure, while sampling in the physical structure reveals parallel, detailed semantic structure. LyberWorld is too abstract to support real spatial navigation: the interface lacks sufficient environmental structure and perceptual cues modeled on the physical world. The Bead system, however, supports navigation in a richer VR environment. Moreover, it explicitly adapts Lynch's five legibility features to enhance virtual wayfinding. The landscape's structure (organizational scheme) and legibility cues support mental mapping for survey knowledge; the system displays cost (distance) and benefit (document) information to support route planning; and the landscape features paths and landmarks to facilitate plan execution. The prototype thus demonstrated a successful adaptation of physical principles for electronic worlds. It also showed the utility of auxiliary information dimensions (e.g., color) that enhance a stable information landscape, for both legibility features and dynamic queries. LyberWorld's techniques seem better suited for use as tools in electronic worlds; Bead's techniques are appropriate for developing electronic environments. Bead's display of usage history captured an aspect of real-world experience, as well as showing a potential social aspect of electronic navigation.

2.5.5 Spaces for Public Information

As discussed above, private information spaces can often be complex and difficult to navigate. In public spaces, e.g., the WWW, these problems can be compounded. At least three solutions exist for this problem: designing better structures, adding navigable superstructures, and designing better visualizations. These three solutions can be combined. The addition of navigable superstructures has been discussed above in connection with effective view navigation (Furnas, 1997) and a Web visualizer with graphical query support (Mendelzon, 1996). This section will discuss additional research that enhances complex information spaces with navigable, spatial presentations.

Adopting a hybrid approach for physical and semantic structures, Dieberger (1995, 1996) developed a prototype that combines textual VR with a Web browser. He proposes that spatial UI metaphors facilitate user navigation in hypertext, by introducing location, distance, and direction. Dieberger's Juggler system is a "multi-user dungeon, object-oriented" (MOO). Derived from online text adventure games, MOOs feature information objects, avatars, and spatial and architectural metaphors. Users can form accurate mental representations from textual descriptions in MOOs. In Juggler, objects and locations can be associated with Web documents. MOO navigation by spatial metaphor causes documents to be loaded in a slave Web browser. This process thus overlays a navigable superstructure on complex Web structures (c.f. Furnas, 1997). The imagined environment is organized as a semi-regular grid. MOO locations can have three types of exit: *directional*, *non-directional*, and *special* (a "magic" feature that functions as a hypertext link). Combinations of such exits can support three types of navigational topology: *Euclidean* (physical), E*uclidean with exceptions* (parallel physical and semantic), and *non-spatial* (semantic).

Juggler uses a Euclidean topology with exceptions (c.f. Kaplan and Moulthrop, 1994). The new Euclidean structures offer a reduced, more navigable version of Web structure. For usability, spatial structures require appropriate navigational metaphors. Small-scale metaphors include transportation modes; large-scale metaphors include city structures. Since MOOs support multi-user communication and interaction, Juggler also offers social navigation. Although imaginatively rich, the cognitive indirectness of Juggler's approach to spatial structure is problematic. The growing popularity of graphical, virtual communities suggests that users may prefer more cognitive directness.

The central distinction between public and private information spaces is explored in two papers by Waterworth (1994, 1996). Waterworth notes benefits of innate human skills in navigation and manipulation for spatial UIs. Such UIs follow HCI'S historical trend towards increasing use of these skills. In each proposal, Waterworth presents a three-dimensional spatial model for representing structured information to support user exploration. In presenting general design ideas, Waterworth allows for alternative realization at three architectural levels:

- *structure* - the underlying information organization
- *world model* - the interaction model or UI metaphor
- *user view* - the customized presentation for a particular audience or task

Waterworth's (1994) proposal for Information Islands has two components -- a hierarchically-structured public world and a mobile, private environment. As shown in Figure 2.4, the public world represents a well-ordered domain of information, services, and applications. Its components range from macroscopic to microscopic in scale:

1. *Archipelago* -- a top-level classification of related entities; major classes of service or application
2. *Island* -- the basic semantic unit of the world; a service subclass
3. *Village* -- a cluster of buildings
4. *Building* -- a set of information sources or services with a topic or application focus; includes an interactive directory and an information counter (public agent)
5. *Floor* -- a set of related services; features a lobby with an interactive directory

Figure 2.4. An overview of Information Islands (Waterworth, 1996).

The private world is a personal tool for navigating and comprehending the public space. This tool is modeled as a vehicle supporting 3D movement. The vehicle offers customized private views of the information world. In addition, the vehicle carries a private workspace that collects useful information, services, and applications.

Waterworth (1996) developed two parallel design sketches of spatial worlds for the WWW. "Personal spaces" are proposed to solve Web browsing problems such as multi-threaded navigation, bi-directional history navigation, and users' needs "to casually organize, reorganize, filter and communicate information." Moreover, such information spaces could help to solve explorers' three main problems:

- making sense of encounters
- finding interesting items
- communicating about discoveries

Waterworth's private design is called StackSpace. It proposes a hierarchical model containing slices (WWW pages), stacks, layers, and spaces. Users can move and label items, which decay with time towards the rear of the model. The viewpoint can be moved in three dimensions for effective overview and investigation of information. The related public design, InfraSpace, spatializes an existing Web corpus. In thought experiments, Waterworth concludes that the proposed tools and environments are more effective for information explorers (end-users) than for information providers (services). StackSpace's simple, flexible structures serve private spaces well, but might not deliver complex, highly-structured information as effectively. The author generally recommends the design of intuitively intelligible public spaces, to which users can bring familiar tools and environments.

Juggler and StackSpace both feature spatial worlds for the WWW; they use parallel physical and semantic structures. Information Islands feature a spatial world for a public, hierarchical structure. The design places physical and semantic structures in correspondence at the macroscopic level, and it allows access at the microscopic physical level to parallel semantic structures. Because of its potential for metaphors, including legibility features, Juggler might support effective navigation. It imposes a cognitive burden for wayfinding through the absence of perceptual input, though; locomotion is comparably indirect. StackSpace features relatively abstract legibility features and organizational schemes, which are emergent and evolutionary. Wayfinding is thus partially supported by this design. Information Islands offers good potential for wayfinding support through a graphical, urban metaphor. Legibility features would be essential for a successful realization of this design. In addition, Information Islands use a rich hierarchical metaphor, which follows the structure of mental maps (Chase, 1983; Stevens and Coupe, 1973) and route planning (Passini, 1984), thereby facilitating wayfinding. In all three designs, display information is primarily navigational, since the destinational data exist in a parallel structure. Each design proposes a different distinction between public and private information spaces. The navigational consequences of this distinction require further research. Waterworth's navigational vehicle, in particular, seems a generalizable technique that could be integrated with other electronic environments and tools. The vehicle occupies an intermediate role between environment and tool, being in some senses both. Finally, each world design proposes a usable electronic environment, within which other tools might be used.

2.5.6 Visualizing Hypermedia Structure

Large knowledge or document bases often don't support views of relationships between objects. Users thus have difficulty in forming an overview of the information, as well as establishing connections between objects. Two solutions to this problem have been discussed above. In IR visualization research, an information landscape is synthesized on the basis of automated textual analysis. In prototypical spatial worlds for Web exploration, the user constructs a spatial world with emergent structure. A third approach visualizes knowledge or document relationships explicitly. This approach visualizes the physical structure of linked information items, rather than their logical structure. Hypermedia structure visualizers face particular design challenges. First, unlike IR visualizations or user-structured spatial worlds, hypermedia visualizers have fixed inter-document links; creating a spatial landscape for navigation is not possible. Second, unlike tree visualizers, hypermedia visualizers must handle general graphs; complex topologies must be supported. Effective visual organization thus plays a central role.

The seminal work on hypermedia structure visualization is the SemNet project (Fairchild, Poltrock, & Furnas, 1988). For understanding large knowledge bases, the authors hypothesize, a user must recognize three things:

- the identities of individual elements in the knowledge base
- the relative position of an element within a network context
- explicit relationships between elements

Accordingly, SemNet knowledge bases are presented as 3D directed graphs, which reveal relationships between symbolic entities. In order to exploit human skills in visual pattern recognition and 3D spatial navigation, a knowledge base is explored by 3D navigation and manipulation. To reveal the structure of a knowledge base, spatial layout is essential. Three solutions were tried:

1. The properties of data elements are mapped to graphical locations, thereby matching the user's conceptual model.
2. Inter-element connectivity determines display adjacency, through multi-dimensional scaling, or centroid or annealing heuristics. These techniques reveal subsets and foreground well-connected objects.
3. Users assign positions on the basis of extra-systemic information, as in daily life (n.b., Waterworth, 1996).

To reduce information display for both graphical computing efficiency and human comprehension, SemNet used three types of fisheye view: hierarchical clustering, 3D point perspective, and gradient sampling density. Navigational problems were primarily those of recognizing and controlling locations. For recognizing locations, real-world tools provide useful models – maps, landmarks, and paths. The authors state that "the single most important feature of the interface is to make the user experience a *real*, three-dimensional space" (Fairchild et al., 1988). The quality of the 3D imagery is essential (as in XEROX PARC's Information Visualizer, Card et al., 1991). For controlling locations, SemNet offered five techniques:

- relative movement - can be confusing and slow
- absolute movement with a map and pointing - faster, but inaccurate and conceptually harder
- teleportation (a history list) - works well with other techniques
- hyperspace movement along inter-object links - useful in clustered environments
- moving the space itself - hasn't been systematically evaluated.

In later research, Ingram and Benford (1995) have adapted Lynch's (1960) legibility research to enhance existing visualizations. They sought to make a hypermedia visualization easier to learn and to navigate. Like city residents, users of persistent, stable, and reusable visualizations can benefit from repeated exposure to good designs, which help them to construct cognitive maps. To create good designs, Ingram and Benford developed a system to enhance visualizations automatically, by using existing database and visualization information. Their prototype system, LEADS (Legibility for Abstract Data Spaces), places visual features in an order-dependent way:

1. Districts are found by cluster analysis.
2. Edges are created quickly by drawing a line between the nearest neighbors in adjacent clusters.
3. Landmarks are determined by cluster centroid triangulation, a balanced and data-responsive method.
4. Nodes can first be chosen as gateways between districts and as cluster centroids; paths can join such nodes first in nearest-neighbor pairs, and then in a minimum spanning tree. Better-chosen nodes and paths would require system usage information, such as access frequency.

LEADS has been applied to three visualization systems, including a multi-user, spring-force model, 3D tool for arbitrary networks. In this case, several hundred WWW nodes were handled with six degree-of-freedom navigation, node inspection and/or Web browser use, and manual node repositioning. Ingram and Benford report generally positive user experiences. Future goals include better 3D path navigation and improved general navigational tools. Related later work is reported in Chalmers et al., 1996.

Another spring-force model 3D network visualization, the Narcissus system, has been developed for software engineering and Web applications. This system combines a self-organizing system with VR. User navigation and comprehension are facilitated by system-controlled object organization, which generally stabilizes quickly. Both static and dynamic equilibrium reveal useful emergent structure. Node clusters are enclosed by translucent surfaces, for aggregation with visual distance. Researchers are currently experimenting with 3D texture mapping and icons, to add semantic information to cluster representations. A prototype implementation integrates Narcissus with the Mosaic Web browser.

The File System Navigator (FSN), a hypermedia structure visualization, is distributed with Silicon Graphics (SGI) computers. This application uses a landscape metaphor to display a UNIX file system for visual navigation and manipulation. Directories appear as rectangular cities, topped by blocks representing files. The directories are connected by linear paths on a green plain. The file system's root is initially positioned at the front of the scene, while directories branch and recede towards the horizon. The user's viewpoint can be moved in three dimensions; a fisheye effect is achieved through perspective geometry. File attributes are shown by graphical attributes, and traditional file management features are supported.

Modeled on FSN, a hypermedia structure visualizer has been implemented in the Harmony client for the Hyper-G Internet hypermedia system (Andrews, 1995; Andrews, Kappe, & Maurer, 1995). The Harmony browser offers several tightly-coupled, 2D and 3D visualizations and navigational tools:

- The *Session Manager* supports general features such as collection navigation, search facilities, and administrative functions. It always displays the navigational path to the current document.
- The *Local Map* displays a dynamic neighborhood map for the currently browsed document. The map displays hyper-links, collection membership, or other attributes.
- The *VRweb 3D scene viewer* shows models in VRML and other file formats. In this visualizer, users navigate and manipulate objects in three dimensions, activating embedded hyper-links as needed.
- The *Information Landscape*, an adaptation of the FSN, is tightly coupled with the 2D Session Manager.

Researchers in the Graphics, Visualization, and Usability Laboratory at the Georgia Institute of Technology (GIT) developed the Navigational View Builder (NVB), a navigational visualization of complex hypermedia networks such as the WWW (Mukherjea et al., 1995; Mukherjea and Foley, 1995). NVB uses four strategies:

- binding data attributes to graphical ones
- content- and structure-based clustering for abstraction
- content- and structure-based filtering for information reduction
- hierarchization for effective visualization.

Because traditional network overview diagrams don't scale well, the NVB features a set of hierarchical sub-views, each with a perspective on the information domain. Analysis structures the data into a pre-tree, a loose hierarchy of hierarchies. The results are displayed using hierarchy visualizers. To generate views, the NVB uses both content and structural analysis. (Web meta-data is currently inadequate for this process.) Although the process is normally automatic, the user can guide the translation of network to a hierarchy, and the visualization of the resulting hierarchy.

At MIT's Media Lab, Rennison (1994) has developed the Galaxy of News visualization system. A prototype USENET news system, Galaxy drew on fisheye views, cone trees, and Pad to support user navigation in an abstract, multidimensional information space. The author's design goals included the following:

- effective exploration and browsing of large news databases
- combining filtering with browsing
- assisting comprehension of relationships between articles
- dynamic content-based hierarchical structuring.

To use the system, the user navigates visually from the general to the specific - from general keywords, through specific keywords, through article headlines, to article text. Mouse selection of an item centers it in a fisheye view and zooms the user's viewpoint towards it. Visible lines indicate hierarchical relationships. The system model includes four layers: the news base, the information relationships, the specification for spatial construction, and the specification for temporal and behavior interaction. The visual layout depends entirely on current news articles. The prototype investigates several issues:

- pyramidal structuring and presentation (for details with context)
- content-based clustering
- abstract 3D spaces
- semantic zooming (for searching or filtering) and panning (for browsing)
- graphical animation
- dynamic navigational cues and layouts - font and article size, text and link transparency, spatial layout, and text and link color.

Emphasizing query features as well as visualization, research in the database lab at the University of Toronto (UT) produced a set of tools (Mendelzon, 1996). Combining a sort of dynamic query with hypermedia structure visualization, one tool displayed 2D network graphs using a variety of algorithms. Small histograms below a graph reveal additional attributes of the data. Users can refine the data view dynamically by manipulating the graph and the histograms. Adapting database techniques for the Web, the group has also developed a system for WWW visualization and querying, on the basis of both structure (link relationships) and content (document title, etc.). The system allows the user to navigate the Web as a 2D graph. The system also supported a sort of semantic zooming through sub-graph nesting to arbitrary depth. The system featured a visual query language, which used the same elements as the visualization: "queries are annotated (hy)graph patterns whose matches on the database can be displayed in different ways ." A slave Web document browser allowed the user to move between structural and document views. (A slave browser is also used in Dieberger, 1996; Waterworth, 1996.)

The hypermedia viewers discussed in this section all handle graphical structures, whether tree (FSN), directed graph (SemNet and Galaxy of News), or general graph (LEADS, Narcissus, Hyper-G, Navigational View Builder, and UT tools). In most cases, physical structure corresponds to semantic structure; in SemNet and the Navigational View Builder, physical structure is derived from semantic structure through analysis. With designs incorporating legibility elements, FSN, Hyper-G, and LEADS support good wayfinding. The FSN, in particular, benefits from a simple tree structure. SemNet investigates several design parameters affecting navigation, including layout and navigational mechanisms. Although anecdotally difficult to use, SemNet demonstrates the value of good spatial structure and cues for navigation. Narcissus' self-organization may generate informative layouts, but the technique may be too dynamic for good wayfinding. This prototype's semantic zooming may somewhat mitigate the effects of layout instability. Although it lacks good spatial cues to support wayfinding, the Navigational View Builder could facilitate mental mapping through inherent hierarchical structure. All of the visualizers reviewed in this section offer data that is simultaneously navigational and destinational. Relative to textual document lists, these viewers generally demonstrate the value of graphical views for overview and exploration (as in IR visualizations). Relative to 2D structure views, these viewers also demonstrate the value of 3D graphics for displaying large amounts of information without occlusion (as in XEROX PARC's Information Visualizer). Relative to standalone hypermedia document viewers, these visualizers show the navigational value of displaying structural context.

2.5.7 Discussion

A diversity of visualization techniques has been reviewed above. They generally place physical and semantic structures in correspondence, although several techniques structure them in parallel. Both highly and loosely structured data have been visualized. Loose structures, such as those in the WWW and large document corpora, are difficult to visualize and manipulate. For complex information, several techniques have been developed for balancing local detail with global context. In most cases, visualization techniques are specialized for specific information structures; of these, hierarchical structures appear often in both hypermedia and spatial worlds. In general, layout organization can be mapped directly from the underlying data domain, derived from the data domain by the system, or emerge through user-system interaction. Most visualizers contain features that conform to wayfinding recommendations. These features may be inherent, derived, enhanced; some explicit adaptation of real-world principles has been successful. Dynamic queries seem particularly useful for facilitating the addition of attributive information to mental maps. Although very dynamic environments may hinder wayfinding, such environments may be useful for other reasons. Animated transitions during locomotion may improve wayfinding, especially for relative movement. Nevertheless, a key visualization challenge is that of maintaining user orientation in deeply nested navigational contexts. For WWW visualization, several hypermedia or spatial visualizations have been developed. These visualizations accord with Furnas' (1997) recommendations for navigable superstructures in complex structures.

A key design goal of information visualization is to support the user in the processes of pattern recognition and aggregation/abstraction. Related system problems include the display of large quantities of information while retaining legibility. For navigation, physical realism through UI metaphor and computer graphics is essential. In addition, users require system support for visualizing tradeoffs between navigational cost and informational benefit. There exists some tension between the predominant approach that seeks to off-load cognition onto perception (e.g., Card et al., 1991), and the opposite approach that seeks to generate strong cognition through basic perceptual cues (e.g., Dieberger, 1996). For the future, many of the reviewed techniques would benefit from validation and/or refinement by behavioral testing. Extrapolation of these techniques to immersive VR would also be instructive. Research opportunities exist for integrating different combinations of the reviewed techniques, particularly those that can be paired as environment and tool. Ultimately, effective computational and interaction techniques are both required for effective navigation in electronic worlds.

As a general summary, Card et al. (1999) offer the following list of key issues for future research in information visualization:

1. New metaphors and visualizations
2. Bringing science to the craft
3. The visualization of cyberspace
4. Collaborative visualizations
5. A characterization of information visualization down to the (data) operator level
6. The perceptual analysis of dynamic information displays
7. Advances in the science of dynamic spatial cognition
8. A theory of knowledge crystallization (i.e., reaching the most compact description of a data set relative to a task)

In terms of the themes developed in this chapter, issues 1, 2, 3, and 7 seem particularly promising.

2.6 Summary

This chapter has reviewed research on navigation in electronic worlds. Three major research areas have been discussed: structuring, exploration, and visualization. Due to problems of size and complexity, electronic worlds are not a panacea for problems of information access, management, and communication, but techniques and tools from research on exploration, structuring, and visualization can enhance navigation in electronic worlds. Good navigation can improve the user's understanding and memory of information; their ability to make informed cost-benefit tradeoffs; and their querying or browsing effectiveness. Ultimately, navigational techniques and tools may prove equally important for task performance as individual user differences. In particular, good internal (mental) and external (physical or electronic) maps are essential for effective wayfinding. Fortunately, principles from social science - especially mental maps and spatial knowledge - and urban design - especially legibility features and wayfinding design - can often be adapted for electronic worlds.

In general, hierarchical structures are often intuitive and effective as a basis for hypermedia and spatial worlds. But general graphs (e.g., the WWW) and unstructured data (e.g., document corpora) are also important. Complex or poorly navigable structures can be improved by adding a navigable superstructure. In any human-computer interaction, there are inherent gaps between a user's cognition, the semantics of the electronic world, and the physical representation of this world. To support navigation, these gaps can be reduced in several ways: UI architectures oriented to user psychology; appropriate high-level UI metaphors; and well-designed relationships between physical and semantic structures. These gaps can be also reduced by careful consideration of user and task in the design of electronic worlds and tools.

There exist many types of electronic environments, navigational tools, navigational tasks, and user strategies. Consequently, there are many opportunities for research and design in hypermedia and virtual worlds. A short summary includes the following opportunities:

1. a taxonomy of tasks, environments, and strategies
2. a taxonomy of navigational tools and designs
3. the development and application of structuring and visualization techniques for electronic worlds
4. the role of non-visual sensory modalities, e.g., audio or kinesthetic cues
5. the nature and role of spatial and temporal cues
6. implications for social navigation, CSCW, and virtual communities

2.7 Conclusions

In seeking the intersection of the above list of opportunities with that proposed by Card et al. (1999), the third item above comes to the fore: the development and application of structuring and visualization techniques for electronic worlds. This is one of the key themes of the current research. The following items then specify potential issues to pursue within this broad research area: new metaphors and visualizations, bringing science to the craft, the visualization of cyberspace, and advances in the science of dynamic spatial cognition. In considering these issues, the activity of user navigation remains central.

This is another key theme of the current research. For such navigation, the mapping between structure and semantics is regarded as one of the most challenging (and hence interesting) issues for contemporary visual UI designers (Jul and Furnas, 1997). For research on these issues, the WWW is a particularly promising information domain, on account of its growing size and importance.

The preceding considerations led the thesis research to focus on the following objectives, as noted in the introduction to this document:

- validation of visualization designs using primarily an landscape/urban metaphor
- investigation of the relationships between hierarchical data structure and UI strategies
- development of an appropriate search task methodology for this domain
- identification of objective versus subjective factors of user performance
- consideration of differences in spatial ability, particularly with regard to user navigation
- consideration of individual differences in structure-learning ability, particularly with regard to hierarchical data

Specific experimental factors to be considered included the following: the style of visual object cues; dimensionality (3D versus 1+D), and the presence or absence of visual objects; the type of visual structure; and the choice of user viewpoint. For this research, desktop VR provided a new and flexible operating environment, while the Web offered a data source with both research and general interest.

In the following chapters of this thesis, a series of studies will be presented that pursue the research ideas discussed above. As research in this field is relatively new, no general theoretical framework is available. By necessity, then, this research is exploratory, proceeding in a series of iterations guided by expectations, rather than by formal hypotheses. One of the contributions of this research is thus the methodology itself.

Chapter 3

Study 1 - Day, Dusk, and Night Worlds

3.1 Introduction

A large design space exists for building a VR landscape to visualize information. In particular, a discrepancy between semantic and spatial structure must be reconciled (Dillon et al., 1993). “Users feel familiar and comfortable with systems based on . . . spatial models. From the point of view of a designer, it is a natural choice to build an electronic environment similar to the real world, so that users can easily adopt and transform their interactive behavior, styles, and patterns, from the physical world into virtual ones.” (Chen, 2000). At the same time, such models do not have unlimited descriptive or explanatory value for visualization, which implies the need to seek the correct balance of semantic and spatial structure, which was one of the design goals of the present research.

Our initial hypothesis was that varying the strength of spatial cueing in world designs would significantly affect search performance and environmental perception. More specifically, we expected that more spatial cueing would result in better understanding of the environment’s information items and spatial structure. Accordingly, a set of three virtual worlds was designed on the basis of discussions at the University of Toronto. The goal was to support a study along the lines discussed above, using worlds with different levels of visual intensity for virtual objects, but consistent textual labels and spatial structure.

3.2 Methodology

Twelve participants were recruited for this study at Umeå University. They were graduate students and instructors in the Departments of Informatics, Linguistics, and Information Studies. (The age range of participants was approximately 20-50, and one third of them were female. All participants were fluent in English, as required by the Swedish educational system and confirmed in conversation with the experimenter.) The primary criterion for selection was Web browsing experience. This criterion automatically included three factors – computer experience; proficiency with a graphical user interface (windows, icons, mouse, pointer, and scroll bars); and familiarity with large, hierarchical information structures.

The study was a one-factor within-subjects design with three levels of the factor. This factor was world design: three points on a hypothesized continuum between text-and object-based representations of information structure. Exposure order was fully counterbalanced, with the six possible combinations of world design and exposure order equally represented in the experimental trials. (Details of the experimental design can be found in Appendix A.)

3.2.1 Apparatus

3.2.1.1 Description

During the experiment, participants used an SGI Onyx 2 workstation to explore three desktop virtual worlds. The workstation was equipped with normal input/output devices, including keyboard, speakers, and mouse. (The monitor measured 20” diagonally.) The virtual worlds were viewed in the CosmoPlayer plug-in from SGI for Web browsers. The CosmoPlayer was controlled by moving the mouse on a pad, in conjunction with occasional keyboard modifier keys (e.g., shift, control). By offering users 3D control over point of view and distance, this user interface implicitly allowed users to trade off between detail and overview in their navigational activities.

Each virtual world design applied the idea of Information Islands (Waterworth, 1994) to visualize a filtered subset of a Web index. The three worlds were maximally isomorphic in features (e.g., locations, sizes, and labels). The data set was chosen for interest to participants and the research community. The set included about 1500 items over seven levels of a hierarchy, which allowed for rich detail and computational tractability.

Data were filtered in three ways to fit the landscape/geographical metaphor. First, nodes below the seventh level were pulled up when possible, by recursively deleting parent nodes and promoting their former children one level. Second, leaf nodes above the seventh were pushed down when possible, by recursively inserting new parent nodes above pairs of siblings and demoting the children one level. The inserted parent was named after the two child nodes, with an ampersand separating the halves of the combined name. Finally, remaining leaf nodes above or below the seventh level were deleted.

The first design was an abstract, urban, daylight landscape, with colored objects, directional lighting, and grayscale labels. This design, the *Day World*, had strong color and lighting cues (Figure 3.1). Virtual objects were laid out to maximize imageability by Lynch’s guidelines: islands, cities, neighborhoods, and buildings (“districts”); mountains and rivers (“edges”); rivers, roads, and bridges (“paths”); and geometric objects (“landmarks” or “nodes”). Objects at each level were clustered around landmarks, according to sibling groups in the data hierarchy. Color assignment grouped buildings in neighborhoods with common palettes, while ground and water objects had naturalistic color. Each object had a text label. To avoid information overload, the distance from which a label was visible varied inversely with the label’s depth in the data hierarchy. In general, the best point from which to survey a region was its center.

Figure 3.1. A view of the Day World design.

Part-way between object- and text-based representations of information structure, the *Dusk World* resembled the Day World with significant changes in coloring (Figure 3.2). In this world, objects were desaturated 90% (as in twilight) and semi-transparent. Labels, however, had bright, saturated colors, which were grouped by sibling relationships in the data hierarchy. The design was intended to support shifts between perceptual modes, i.e., object- and text-based.

Figure 3.2. A view of the Dusk World design.

At the textual end of the design continuum was the *Night World* (Figure 3.3), which is essentially the Dusk World without virtual objects. Here, the user could move in an abstract information space, without absolute location or distance. This design was inspired by prototypes such as HotSauce and others (Rennison, 1994). The Night World lacked directional lighting: it showed only brightly colored text on a black background.

In these visualization environments, users could navigate in either discrete or continuous fashion. Like in hypertext, a user could point-and-click to navigate up or down the data hierarchy, and between sibling data nodes. Unlike in hypertext, a user could also navigate continuously between items of interest, even if these items were separated by one or more levels of data hierarchy, or were not located on the same branch of the hierarchy.

Figure 3.3. A view of the Night World design.

3.2.1.2 Development

The experimental VR environments were constructed in VRML 2.0, based on a filtered subset of WWW structure. The VRML was generated by C++ code (approximately 7000 lines), requiring about seven months of development effort. The world-generation algorithm took as input ASCII data in a hierarchical format: processing time is about five seconds per world, and each VRML file is about three megabytes in size. Such desktop VR could be viewed with the CosmoPlayer plug-in for Web browsers. Special input/output devices were not needed. User interactivity was facilitated by using a high frame-rate graphics card.

The experimental environments were developed with the assistance of two industrial designers at the Institute of Design at Umeå University. Initial discussion produced a series of sketches, which were evaluated for expected usability and appropriateness for research goals. Once a consensus was reached on basic designs, the designers specified a set of virtual building forms and a palette of approximately 300 colors. After the first study, design iteration was driven by usability feedback from study participants.

3.2.2 Procedure

Each experimental session lasted approximately two hours per participant, including a rest break. The session began with an introduction to the experiment, general research context, and lab equipment to be used. This introduction took approximately five minutes.

The second stage of the session was devoted to explanation of experimental procedures, as well as training and practice in the user interface for the experiment. First, participants received a detailed explanation of experimental procedures, which took approximately five minutes. Next, participants learned the experimental UI. The training included use of the VRML browser's "Home" button, which was always available, in light of research findings that users of large information environments tend to exhibit strong homing behavior (Catledge and Pitkow, 1995; Tauscher and Greenberg, 1997). Participants were instructed not to use the VRML browser's history mechanism (for backtracking), in order to focus their attention more directly on the navigational environment. The first five minutes of this learning time were devoted to navigation techniques and practice. Participants were then allowed 10 minutes to freely explore a virtual world that represented an Aztec temple compound. The goal of the introduction, training, and practice was to teach each participant the necessary techniques for performing the experimental tasks successfully, as well as achieving psychological comfort. The participant was advised that questions about experimental procedures, UI techniques, and document language would be permitted and answered, if possible, at any time during the experimental session.

During the third stage of each session, each participant engaged in a sort of scavenger hunt (i.e., an ordered search for a series of targets in a large environment) for 20 minutes. A search task was chosen for the study because it could be evaluated unambiguously, and because it related to tasks that actual users need to perform. The investigator began each hunt by locating the participant in the center of a visualization environment (Day, Dusk, or Night World design), and then handed the participant a target description on a paper card. (For the experiment, three sets with 10 targets each had been selected at random from the visualized data set). Each card described an informational context (e.g., "Broadway Musicals"), and then presented a specifically-named target (e.g., "Fiddler on the Roof"). Each participant was allowed five minutes to find each target. When a target was found, the participant was to mouse click on the target to play a recorded sound. If a target was not found after approximately five minutes, the investigator gave the participant a new target card, and the cycle was repeated. If a target proved difficult, a participant was permitted to skip it and proceed to the next one. The participant was requested to find as many targets as possible in the time allowed. Each participant had the option to return to the center of the model world after each target card (indeed at any time), or to remain in their current location. After 20 minutes, the hunt was terminated. The participant was then interviewed briefly about the experience of the virtual world. Audio comments were recorded. Participants were also asked to sketch the virtual world on white paper with colored pens, for which they were allowed 10 minutes. The third stage of the session was followed by an optional five-minute break, as were the third and fourth stages described below.

The fourth stage of the session used a different virtual world design than the third stage, but was otherwise similar. The fifth stage of the session used a different virtual world design than the third and fourth stages, but was otherwise similar.

The experimental session concluded with a two-page questionnaire. In this questionnaire, participants were asked to subjectively rate the model worlds for perceived size, block duration, sense of presence, ease of use, and enjoyment. The questionnaire also asked about age, occupation, languages, computer usage, computer/video game usage, and Yahoo Web index usage. Finally, the questionnaire solicited general comments. Upon completion of this questionnaire, each participant was paid $20 for participation in the experiment. The participant was also debriefed by the experimenter about the purpose of the study.

A complete experimental script and related questionnaires can be found in Appendix A.

3.2.3 Measures

The following measures were used in analyzing the results of Study 1.

- **Independent Variables**
  - *BLOCK*: block of trials (1, 2)
  - *DESIGN*: virtual world design (Day, Dusk, or Night World)

- **Dependent Variables**
  - *TARGETS*: number of search targets found, i.e., successful trials (0 - 9)
  - *DURATION*: rated block duration (1=short to 5=long)
  - *SIZE*: rated world size (1= small to 5=large)
  - *PRESENCE*: self-reported “sense of presence” (1= low to 5=high)
  - *EASE*: self-reported “ease of use” (1= hard to 5=easy)
  - *ENJOYMENT*: self-reported “enjoyment” (1= low to 5=high)

3.3 Results

This study sought to identify factors, relationships, and effects that held across the three visualization environments, in order (1) to iterate the designs with regard to usability and (2) to understand the key factors of user response in this research domain. For this reason, study results will generally be analyzed across the experimental conditions, rather than within specific conditions.

3.3.1 Description

	N	Minimum	Maximum	Mean	Std. Deviation
TARGETS	36	0	9	4.28	2.46
SIZE	36	1	5	3.22	1.12
DURATION	36	1	5	3.11	1.06
PRESENCE	36	1	5	3.00	1.20
EASE	36	1	5	2.81	1.14
ENJOYMENT	36	1	5	2.94	1.19

Table 3.1. Descriptive statistics for the Study 1 measures.

The subjective measures generally clustered about the central value of the five-point scale. In terms of self-reported measures, participants rated the virtual worlds overall as having a moderate sense of presence, ease of use, efficiency, and enjoyment. There was a relatively large number (14) of “Hard” or “Very hard” ratings for the ease of use category. (From an experimental point of view, it is beneficial to have tasks that are neither too easy nor too hard, in order to generate maximal differentiation of dependent measures. With a central mean and well-distributed user ratings for ease of use, the present study seems to have achieved good differentiation.)

The distribution for number of targets found was positively skewed, with the majority of observations in the lower levels of performance. Estimates of block duration were generally clustered around the “Medium” rating.

Participants’ world sketches were suggestive for exploratory research, but difficult to analyze quantitatively. For this reason, they were not used further in the thesis research. (Two sample drawings are located in Appendix C.)

Detailed descriptive statistics for the experimental measures can be found in Appendix B.

3.3.2 Underlying Factors of Performance

In order to examine the relationships between response measures, a matrix of pairwise correlations was derived (Table 3.2). Three of the subjective measures – sense of presence, ease of use, and overall enjoyment – all correlated significantly with each other ($p < .05$). Number of targets found correlated significantly with two of the subjective measures ($p < .05$) – ease of use and overall enjoyment. In addition, rated duration correlated significantly ($p < .05$) with rated size and ease of use (negatively). There were no other significant correlations.

	TARGETS	SIZE	DURATION	PRESENCE	EASE	ENJOYMENT
TARGETS	1.000					
SIZE	-.013	1.000				
DURATION	-.220	**.409**	1.000			
PRESENCE	.126	-.064	.067	1.000		
EASE	**.478**	-.143	**-.288**	**.481**	1.000	
ENJOYMENT	**.405**	-.225	-.265	**.420**	**.683**	1.000

Table 3.2. Correlations for the Study 1 measures (Pearson, one-tailed, N = 36), bold font shows $p < .05$.

The factor structure underlying the performance-related and five self-reported variables was uncovered with factor analysis. Using the scree plot (Figure 3.4), one factor was selected that accounted for approximately 42% of experimental variance. I will first examine the effect of different experimental conditions on this discovered factor, and then examine the effects for each of the original dependent variables.

Factor analysis used “varimax” rotation of a principal factor solution. To reach this solution, a cloud of data points was projected onto a set of orthogonal axes that accounted for the maximum variance within the data. The axes (factors) were then rotated so as to maximize the variance in loadings (hence, "varimax") on the first factor, in order to increase this factor’s interpretability. The goal was to create a reduced set of explanatory factors, so as to clarify the main constructs underlying the set of dependent variables used in the study. This procedure was also used in the remaining factor analyses reported in this thesis.

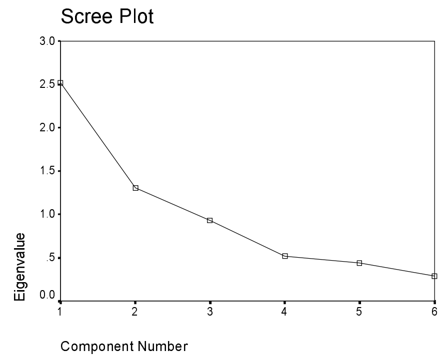

Figure 3.4. Extracted factors vs. variance explained in Study 1.

Varimax rotation was used to enhance the internal probability of the solution, and the resulting factor loading for the rotated solution is shown in Table 3.3.

	Component
TARGETS	**.60**
SIZE	-.05
DURATION	-.15
PRESENCE	**.71**
EASE	**.87**
ENJOYMENT	**.82**

Table 3.3. Rotated component matrix for Study 1, one-factor solution (bold font shows factor loadings above .5 or below -.5).

Factor 1 is composed primarily of “feeling” (attitudinal) and “doing” (behavioral) measures – targets found, sense of presence, ease of use, and overall enjoyment - as indicated by the high loading for these measures varying in absolute value between .60 and .87.

In order to examine the effect of world design on the underlying factor, I performed a one-way ANOVA using design as a within-subjects measure. There was a significant effect of design on Factor 1 ($F_{2,22}$ = 7.65, p = .003). Participants generally had a more positive response to the Day World condition than to the other two conditions, as shown in Figure 3.5. (Error bars in these graphs refer to 95% confidence intervals around the mean.) The source of this effect was examined using repeated contrasts (by the method in the repeated-measures ANOVA feature of SPSS for Windows, version 9.0). This analysis showed a significant difference between the Day and Dusk Worlds ($F_{1,11}$ = 15.55, p = .002), and between the Day and Night Worlds ($F_{1,11}$ = 4.85, p = .050), but no significant difference between the Dusk and Night Worlds.

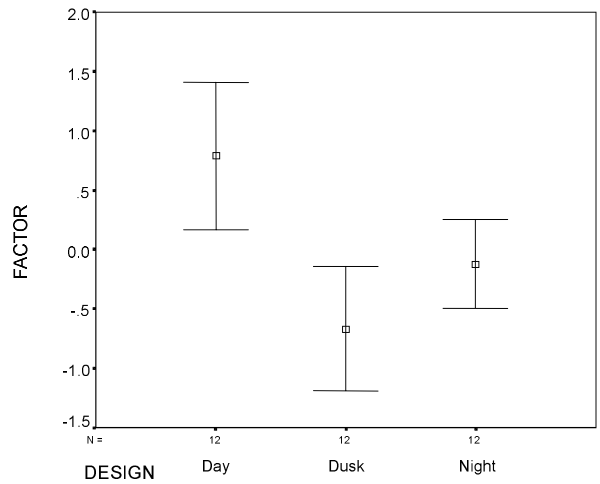

Figure 3.5. Effect of world design on extracted factor 1 in Study 1.

In order to examine the effect of task learning on the underlying factor, I performed a one-way ANOVA using block number as a within-subjects measure. There was no significant effect of task learning on the underlying factor.

3.3.3 Effects on User Strategy

The next analysis examined the particular experimental measures that describe performance and attitude. This inquiry will be presented as effects of two experimental variables: world design and task learning. Statistical analyses will be used to identify significant interactions and main effects involving these variables, as a step towards constructing a model of user behavior and attitude.

3.3.3.1 Effects of World Design

In order to investigate the effect of world design on dependent measures, I performed a series of one-way ANOVAs using world design. There was a significant effect of world design on three of the subjective measures.

There was a significant effect of world design on sense of presence ($F_{2,22} = 4.86$, $p = .018$). Participants felt more of a sense of presence in the Day World (3.83) than in the Dusk World (2.67) or the Night World (2.5). The source of this effect was examined using repeated contrasts. This analysis showed a significant difference between the Day and Dusk Worlds ($F_{1,11} = 15.40$, $p = .002$), and between the Day and Night Worlds ($F_{1,11} = 6.77$, $p = .025$), but no significant difference between the Dusk and Night Worlds.

There was a significant effect of world design on ease of use ($F_{2,22} = 4.09$, $p = .031$). Participants reported a greater ease of use in the Day World (3.42) than in the Dusk World (2.17) or the Night World (2.83).

The source of this effect was examined using repeated contrasts. This analysis showed a significant difference between the Day and Dusk Worlds ($F_{1,11}$ = 12.69, p = .004), but no other significant differences.

There was a significant effect of world design on overall enjoyment ($F_{2,22}$ = 7.24, p = .004). Participants reported enjoying the Day World (3.92) more than the Dusk World (2.17) or the Night World (2.75). The source of this effect was examined using repeated contrasts. This analysis showed a significant difference between the Day and Dusk Worlds ($F_{1,11}$ = 22.15, p = .001), but no other significant differences.

There were no other significant effects of world design on the experimental measures.

3.3.3.2 Effects of Task Learning

In order to investigate the effects of task learning on experimental measures, I performed a series of one-way ANOVAs using block number as a within-subjects measure. There were no significant effects of task learning on the experimental measures.

3.3.4 Learning Effects within Blocks

In order to examine more closely user strategy within each block of trials, I examined the number of targets found during each two-minute period of a block. There was no significant effect of elapsed time on the number of targets found.

I then defined two new measures, H1 and H2. These measures represented the number of targets found during the first and second half, respectively, of each block. I then performed a one-way ANOVA using block half as a within-subjects measure. There was a significant (negative) effect of within-block task learning on number of targets found ($F_{1,35}$ = 6.15, p = .018). Participants generally found more targets during the first half-block (2.44) than during the second half-block (1.83).

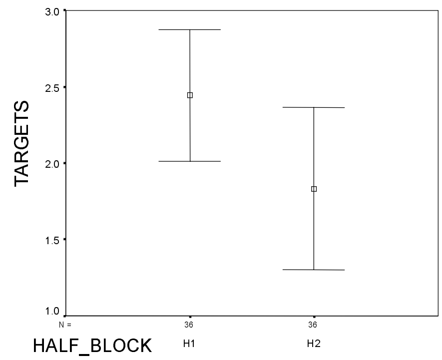

Figure 3.6. Effect of task learning on number of targets found (per half-block).

3.3.5 Experimental Control

In order to examine effects of asymmetric transfer between experimental conditions, I performed a series of two-way ANOVAs using world design and block number as within-subjects measures. There were no significant effects of these interactions, and thus no evidence of asymmetric transfer.

To examine any effects possibly hidden by task learning, I restricted the data to the first block only, and performed a series of one-way ANOVAs using world design as a between-subjects measure. There were no significant effects on the experimental measures.

In order to examine the effects of different target sets, I performed a series of one-way ANOVAs using target set as a between-subjects measure. There were no significant effects of target set.

3.6 Conclusions

The most striking result of Study 1 is the strong user preference for the Day World design. This design was rated superior to the other designs for sense of presence, ease of use and overall enjoyment. Users apparently preferred the relative naturalism of the Day World. Comments by some users suggested that the Night World was often perceived as unfriendly and disorienting, while the Dusk World was often seen as dull.

Another striking result of the experiment was a lack of significant effect of virtual world design on task performance. The constant across world designs was structural isomorphism. As noted in the Methodology section of this chapter, a given object or label had the same Cartesian coordinates and dimensions in each virtual world. This study's results suggest that a key factor in determining user perception for the search task was style of visual representation for information structure.

Study 1 clearly showed a separation between objective and subjective factors of participant response. World design and task learning had no significant effect on performance measures. Yet world design had significant effects on participant attitude and perception.

There appears to have been a fatigue or motivational factor during the experimental blocks, as task performance generally decreased during the second half of each block. This decrease suggests that the experimental tasks were too difficult, the visualization environments insufficiently usable, or the UI training time insufficient. User comments suggested that all of these issues might have played a role. (Later studies addressed these issues, which may help to explain the absence of a negative learning effect in those studies.) In any case, the experimental situation was a complex one, with a number of factors driving performance.

Further research was needed (1) to clarify the psychological issues in user response to experimental conditions, (2) to examine more closely the representation of structure in information visualization, and to further develop the experimental methodology.

Chapter 4

Study 2 - VR vs. Hypertext

4.1 Introduction

Study 1 showed a general user preference for a naturalistic desktop VR environment (Day World) for navigation, relative to two other designs. For Study 2, since hypertext is the form of the world's largest information structure (the WWW), it seemed appropriate to establish a baseline for performance and perception by comparing hypertext with desktop VR. At the same time, Study 2 investigated the effect of varying domain data sets and user interface independently, in order to clarify the relationship between interface design and domain data. For this purpose, a simple hypertext was developed by specialization (sub-classing) in the world-generation software from Study 1 (Figure 4.1).

Figure 4.1. A view of the hypertext design.

4.2 Methodology

Sixteen participants were recruited for this study at Umeå University. They were undergraduate and graduate students, mainly in the Department of Informatics. (The age range of participants was approximately 20-40, and approximately one quarter of them were female. All participants were fluent in English, as required by the Swedish educational system and confirmed in conversation with the experimenter.) The primary criterion for selection was Web browsing experience. This criterion automatically included three factors – computer experience; proficiency with a graphical user interface (windows, icons, mouse, pointer, and scroll bars); and familiarity with large, hierarchical information structure.

The study was a 2 x 2 fully within-subjects design. The first independent variable was world design: desktop VR vs. hypertext. The second independent variable was data set: work vs. leisure subsets of the data in Study 1. A Latin Square design was used to vary exposure order. (Details of the experimental design can be found in Appendix D.)

4.2.1 Apparatus

After Study 1, some usability and performance problems in the prototype design were fixed, largely based on participants' written and oral comments.

During the study, participants used a standard desktop PC to explore two visualization environments. As mentioned, these worlds were the Day World from Study 1 and a hypertext representation of the same data. The PC had a Dell 400 MHz Pentium CPU, 128 MB RAM, a Diamond Fire GL 4000 graphics card, a 21" Trinitron monitor, keyboard, speakers, and mouse. The model worlds were viewed in the CosmoPlayer plug-in from SGI for Web browsers. As in Study 1, the CosmoPlayer was controlled by moving the mouse on a pad, in conjunction with occasional keyboard modifier keys (e.g., shift, control). Mouse and keyboard events were logged to disk files using a Microsoft development utility called Spy++.

4.2.2 Procedure

Each experimental session lasted approximately two hours per participant, including a rest break. The session began with an introduction to the study, general research context, and lab equipment to be used. This introduction took approximately five minutes.

The second stage of the session was devoted to explanation of experimental procedures, as well as training and practice in the user interface for the study. First, participants received a detailed explanation of experimental procedures, which took approximately five minutes. Next, participants learned the experimental UI. The training included use of the VRML browser's "Home" button, which was always available, in light of research findings that users of large information environments tend to exhibit strong homing behavior (Catledge and Pitkow, 1995; Tauscher and Greenberg, 1997). Participants were instructed not to use the VRML browser's history mechanism (for backtracking), in order to focus their attention more directly on the navigational environment. The first five minutes of this learning time were devoted to navigation techniques and practice. Participants were then allowed 10 minutes to freely explore a virtual world that represented an Aztec temple compound. The goal of the introduction, training, and practice was to teach each participant the necessary techniques for performing the experimental tasks successfully, as well as achieving psychological comfort. The participant was advised that questions about experimental procedures, UI techniques, and document language would be permitted and answered, if possible, at any time during the experimental session.

During the third stage of each session, each participant engaged in two scavenger hunts for 15 minutes each, after each of which a questionnaire was administered about the visualized data. The investigator began each hunt by locating the participant in the center of a visualization environment (desktop VR or hypertext visualization; work or leisure data subset), and then handing the participant a target description on a paper card. (For the study, two sets of 25 targets each had been randomly selected for each data set.) Each card described an informational context (e.g., "Broadway Musicals"), and then presented a specifically-named target (e.g., "Fiddler on the Roof"). Participants were allowed approximately five minutes to find each target, after which they were strongly encouraged to proceed to the next target; the participant was requested to find as many targets as possible in the time allowed. When a target was found, the participant was to mouse click on the target to play a recorded sound, and then press the "T" key (to make an entry in the log file). If a target was not found after approximately five minutes, the investigator gave the participant a new target card, and the cycle was repeated. If a target proved difficult, participants were permitted to skip it and proceed to the next one. Each participant had the option to return to the center of the model world after each target card (indeed at any time), or to remain in their current location. After 15 minutes, the hunt was terminated. The participant was then asked to classify five hypothetical targets on a paper classification form offered for this purpose. In creating this form, items were randomly selected from the set of targets not used in the hunt. (A sample form is shown in Appendix D.) The form also asked participants to rate the top-level data categories for perceived size, and to estimate the duration of the scavenger hunt. The participant was allowed as much time as necessary to complete this form. The third stage of the session was followed by an optional five-minute break.

The fourth stage of the session used two visualization environments (hypertext or desktop VR visualization; leisure or work subset) different from those used in the second stage; the fourth stage was otherwise similar, with the addition of a final 1-page questionnaire. In this questionnaire, participants were asked to subjectively rate the model worlds for sense of presence, ease of use, efficiency, and enjoyment. Upon completion of this questionnaire, each participant was paid $20 for participation in the study. The participant was also debriefed by the experimenter about the purpose of the study.

A complete experimental script and related questionnaires can be found in Appendix D.

4.2.3 Measures

The following measures were used in analyzing the results of Study 2.

- **Independent Variables**
  - *BLOCK*: block of trials (1-4)
  - *DESIGN*: virtual world design (hypertext, VR)

- **Dependent Variables**
  - *TARGETS*: number of search targets found, i.e., successful trials (0 - 30)
  - *DURATION*: estimate of block duration (0 – 28 minutes)
  - *SIZE*: estimate of world size (1= small to 5=large)
  - *PRESENCE*: self-reported “sense of presence” (1= low to 5=high)
  - *EASE*: self-reported “ease of use” (1= hard to 5=easy)
  - *EFFICIENCY*: self-reported “efficiency” (1=inefficient to 5=efficient)
  - *ENJOYMENT*: self-reported “enjoyment” (1= low to 5=high)
  - *STRUCTURE TEST*: score on post-block structure test (0 - 1.0)

Participants’ ratings of information category sizes (after each block) and visualization environment sizes (at the end of the session) were inconclusive, and will not be discussed further in this thesis.

For scoring a post-block structure test, the primary criterion for correctness was this: accuracy sufficient to guide correct navigational choice in the block’s virtual world. Each structure test contained five items. Each item’s answer was scored for correctness of the first two nodes named. For each correct first node, matching one level below the root of the data hierarchy, a point was added. For each correct second node, matching two levels below the root of the data hierarchy, a point was added. Participants could receive a maximum score of two per test item, and thus a maximum total score of ten for all five items. Each test's raw score (0-10) was normalized to a maximum value of 1.0.

For example, one structure-learning test contained the following item: “Folk musician Jeff Buckley.” The correct answer was as follows: “Entertainment → Music → Artists → By Genre → Folk.” One participant offered the following answer: “Entertainment → Music → Folk Music → Artists → Jeff Buckley.” For the node “Entertainment,” this answer received a point; for the node “Music,” this answer received a second point.

4.3 Results

This study sought to identify factors, relationships, and effects that held across the two visualization environments, in order (1) to iterate the designs with regard to usability and (2) to understand the key factors of user response in this research domain. For this reason, study results will generally be analyzed across the experimental conditions, rather than within specific conditions.

4.3.1 Description

	N	Minimum	Maximum	Mean	Std. Deviation
TARGETS	60	0	21	7.10	5.35
DURATION	60	4	30	14.18	5.21
SIZE	60	2	5	3.80	.92
PRESENCE	60	1	5	3.17	1.04
EASE	60	1	5	3.00	1.30
EFFICIENCY	60	1	5	3.07	1.27
ENJOYMENT	60	1	5	3.30	1.11
STRUCTURE TEST	60	.00	1.00	.4867	.2587

Table 4.1. Descriptive statistics for the Study 2 measures.

The subjective measures generally clustered about the central value of the five-point scale, with a high of 3.30 for enjoyment. In terms of self-reported measures, participants rated the virtual worlds overall as having a moderate sense of presence, ease of use, efficiency, and overall enjoyment. There was a relatively large number (25) of “Hard” or “Very hard” ratings for the ease of use category.

The distribution for number of targets found was positively skewed, with the majority of observations in the lower levels of performance. The distribution of the post-block structure test was concentrated around the mean, while somewhat positively skewed. Estimates of block duration were generally slightly less the actual duration of 15 minutes, while estimates of world extent clustered around the “Large” rating.

Detailed descriptive statistics for the experimental measures can be found in Appendix E.

4.3.2 Underlying Factors of Performance

In order to examine the relationships between response measures, a matrix of pairwise correlations was derived (Table 4.2). Overall enjoyment correlated significantly ($p < .05$) with three of the subjective measures - sense of presence, ease of use (negatively), and rated efficiency (negatively). Number of targets found correlated significantly ($p < .05$) with estimated duration, ease of use, rated efficiency, and overall enjoyment (negatively). In addition, rated efficiency correlated significantly with estimated duration and ease of use ($p < .05$). There were no other significant correlations.

	TARGETS	DURATION	PRESENCE	EASE	EFFICIENCY	ENJOYMENT
TARGETS	1.000					
DURATION	**.232**	1.000				
PRESENCE	-.133	.046	1.000			
EASE	**.591**	.188	-.199	1.000		
EFFICIENCY	**.603**	**.231**	-.110	**.899**	1.000	
ENJOYMENT	**-.265**	-.027	**.219**	**-.634**	**-.542**	1.000

Table 4.2. Correlations for the Study 2 measures (Pearson, one-tailed, N = 60), bold font shows $p < .05$.

The factor structure underlying the performance-related and five self-reported variables was uncovered with factor analysis. Using the scree plot (Figure 4.2), one factor was selected that accounted for approximately 44% of experimental variance. I will first examine the effect of different experimental conditions on this discovered factor, and then examine the effects for each of the original dependent variables.

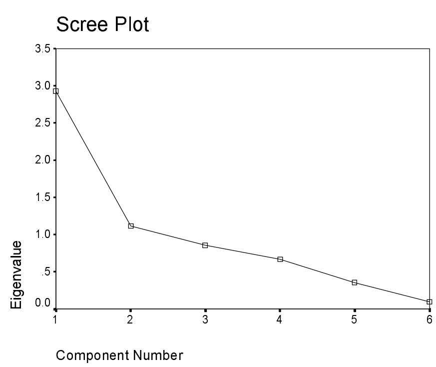

Figure 4.2. Extracted factors vs. variance explained in Study 2.

Varimax rotation was used to enhance the internal probability of the solution, and the resulting factor loadings for the rotated solution are shown in Table 4.3.

	Component
TARGETS	**.74**
DURATION	.33
PRESENCE	-.25
EASE	**.94**
EFFICIENCY	**.92**
ENJOYMENT	**-.69**

Table 4.3. Rotated component matrix for Study 2, one-factor solution (bold font shows factor loadings above .5 or below -.5).

Factor 1 is composed primarily of "feeling" (attitudinal) and "doing" (behavioral) measures – targets found, ease of use, rated efficiency, and overall lack of enjoyment - as indicated by the high loading for these measures varying in absolute value between .69 and .95.

In order to examine the effect of world design on the underlying factor, I performed a one-way ANOVA using design as a within-subjects measure. There was a significant effect of design on Factor 1 ($F_{1,14}$ = 48.91, p = .001). Participants performed better and reported better feelings in the hypertext condition than in the desktop VR condition, as shown in Figure 4.3.

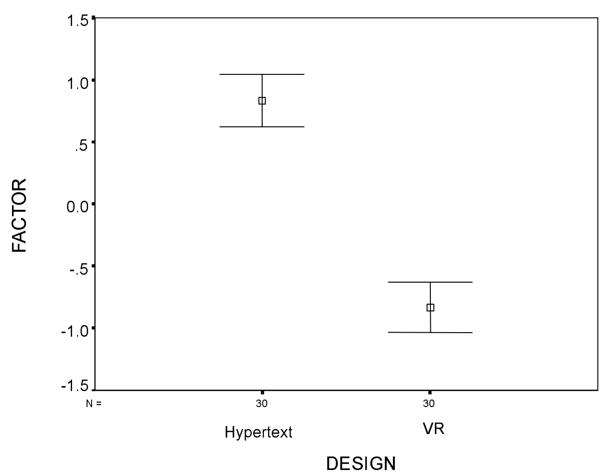

Figure 4.3. Effect of world design on extracted factor 1 in Study 2.

In order to examine the effect of task learning on the underlying factor, I performed a one-way ANOVA using block number as a within-subjects measure. There was no significant effect of task learning on the underlying factor.

4.3.3 Effects on User Strategy

Having considered the underlying factors of performance, it is worth inquiring more deeply into the particular experimental measures that describe performance and attitude. This inquiry will be presented as effects of three experimental variables: world design, task learning, and (domain) structure learning ability. Statistical analyses will be used to identify significant interactions and main effects involving these variables, as a step towards constructing a model of user behavior and attitude.

4.3.3.1 Effects of World Design

In order to investigate the effect of world design on dependent measures, I performed a series of one-way ANOVAs using world design. There was a significant effect of world design on all experimental measures except for sense of presence.

There was a significant effect of world design on number of targets found ($F_{1,14}$ = 79.13, p = .001). Participants found more targets in the hypertext design (11.33) than in the VR design (2.87).

There was a borderline significant effect of world design on estimated block duration ($F_{1,14}$ = 4.55, p = .051). Participants gave longer time estimates for the hypertext design (15.03) than for the VR design (13.33).

There was a significant effect of world design on ease of use ($F_{1,14}$ = 30.0, p = .001). Participants rated the hypertext design (4.0) as easier to use than the VR design (2.0).

There was a significant effect of world design on rated efficiency ($F_{1,14}$ = 21.69, p = .002). Participants rated the hypertext design (4.0) as more efficient than the VR design (2.13).

There was a significant effect of world design on overall enjoyment ($F_{1,14}$ = 6.79, p = .021). Participants rated the hypertext design (2.73) as less enjoyable than the VR design (3.87).

There were no other significant effects of world design.

4.3.3.2 Effects of Task Learning

In order to investigate the effects of task learning on experimental measures, I performed a series of one-way ANOVAs using block number as a within-subjects measure.

There was a significant effect of task learning on estimated duration ($F_{3,42}$ = 3.45, p = .025). Participants in later blocks tended to give larger time estimates than in earlier blocks.

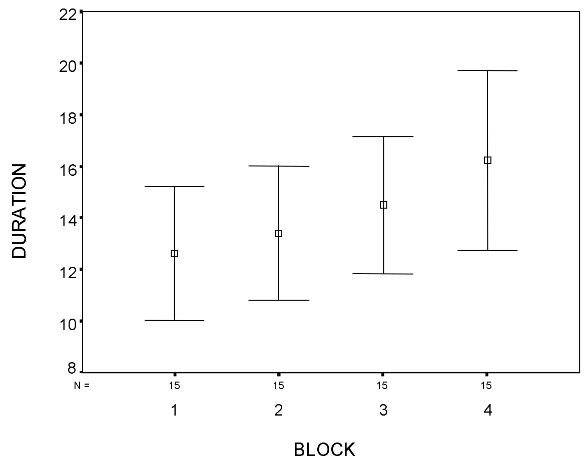

Figure 4.4. Effect of task learning on estimated block duration.

In order to investigate the effects of task learning in specific design conditions on the experimental measures, I performed a series of one-way ANOVAs using block number as a within-subjects measure. In the VR cases only, there was a significant effect of task learning on number of targets found ($F_{3,18}$ = 7.05, p = .002).

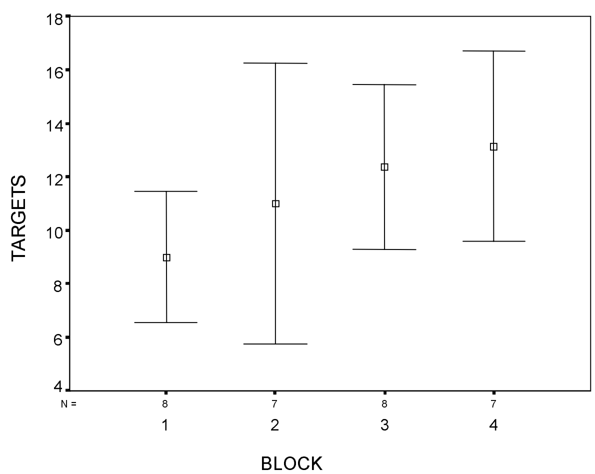

Figure 4.5. Effect of task learning on number of targets found (VR cases only).

There were no other significant effects of task learning.

4.3.3.3 Effects of Structure Learning Ability

In order to investigate the effects of structure learning ability on dependent measures, I averaged each participant's two post-block classification scores, and then aggregated these averages into two groups (high and low). This aggregation created a between-subjects pseudo-factor called STRUCT, which is intended to capture each participant's ability to learn domain information, regardless of effects from virtual world design or task learning. (Speculation suggests that structure-learning ability may be analogous to visual or map memory, but this is a topic for future research.) I then performed a series of one-way ANOVAs using structure learning ability as a between-subjects measure. There were no significant effects of structure learning ability.

4.3.5 Experimental Control

In order to examine effects of asymmetric transfer between experimental conditions, I performed a series of two-way ANOVAs using block number as a between-subjects measure and (1) world design as a within-subjects measure or (2)structure learning ability as a between-subjects measure. There were no significant effects of these interactions, and thus no evidence of asymmetric transfer.

To examine any effects possibly hidden by task learning, I restricted the data to the first block only, and performed a series of one-way ANOVAs using world design and structure learning ability as between-subjects measures. There were significant effects of world design on number of targets found ($F_{1,13}$ = 33.08, p = .001), estimated duration ($F_{1,13}$ = 13.35, p = .003), ease of use ($F_{1,13}$ = 16.18, p = .001), and rated efficiency ($F_{1,13}$ = 15.69, p = .002). As previously discussed, these effects were all significant for overall experimental sessions as well. The significance of these effects for within- and between-subjects analysis tends to confirm their presence in the experimental conditions. There were no other significant effects of block-one analysis.

In order to examine the effects of different data subsets, I performed a series of one-way ANOVAs using data subset as a between-subjects measure. There were no significant effects of data subset.

4.4 Conclusions

In considering results of Study 2, the superiority of hypertext for most measures was evident in number of targets found, ease of use, and rated efficiency. Objective and subjective measures were for the most part in agreement. Hypertext is evidently a more mature technique for information access, as reflected in better software tools and extensive, prior user experience. At the same time participant enjoyment didn’t match the most efficient user interface. In fact, enjoyment was negatively correlated with performance, suggesting greater engagement and motivation in the VR design. Some of this engagement, of course, may stem from novelty of the interface condition. In any case, the experimental results suggest that hypertext’s structural benefits should inform further development in visualization of hierarchical information.

Similarly, the usability advantages of hypertext could be considered potential standards for desktop VR development. Accordingly, participants’ suggestions for improvements of the research prototype, particularly those inspired by comparison with hypertext, had the potential to advance the iterative design process. The resulting prototype could then become a better tool for task completion, as well as a more solid and sensitive platform for future research.

Despite hypertext's advantage on the experimental measures of Study 2, it seems likely that all techniques for presenting hierarchical information would suffer from problems of scale. The current study used approximately 1500 information items; a research opportunity would lie in considering data hierarchies that are substantially larger. For such research, fisheye views (Furnas, 1986) may be a useful concept, which could be applied to many different rendering approaches.

The study's results for time perception are interesting, but inconclusive. According to Waterworth (1983), people may evaluate passing time as shorter when they are engaged in conscious processing of information, rather than in perceptual tasks. In the current study, block duration estimates increased over the course of an experimental session. The cited research suggests the possibility that participants' cognitive processing lessened (shifted towards perception) as they learned the interface and information data set. (In fact, participant performance did improve slightly over the course of an experimental session, though the improvement was statistically significant only in the VR cases.) Similarly, current study participants estimated hypertext block duration as longer than VR block duration; this result may suggest a lighter cognitive load in the hypertext interface. This speculation concurs with the greater sense of involvement (enjoyment) that study participants appear to have felt in the desktop VR environment.

Chapter 5

Study 3 - Naturalistic vs. Efficient Worlds

5.1 Introduction

For search task performance and attitudinal measures, Study 2 showed the advantages of hypertext over naturalistic desktop VR visualization of hierarchical data. For Study 3, it was decided to investigate the differences in visual structure by focusing on the layout algorithm for the visualization model. For this model, some of the structural advantages of hypertext were imposed more firmly on the VR visualization in the form of the CityScape algorithm. An "efficient" Day World was thereby created. (Figure 5.1)

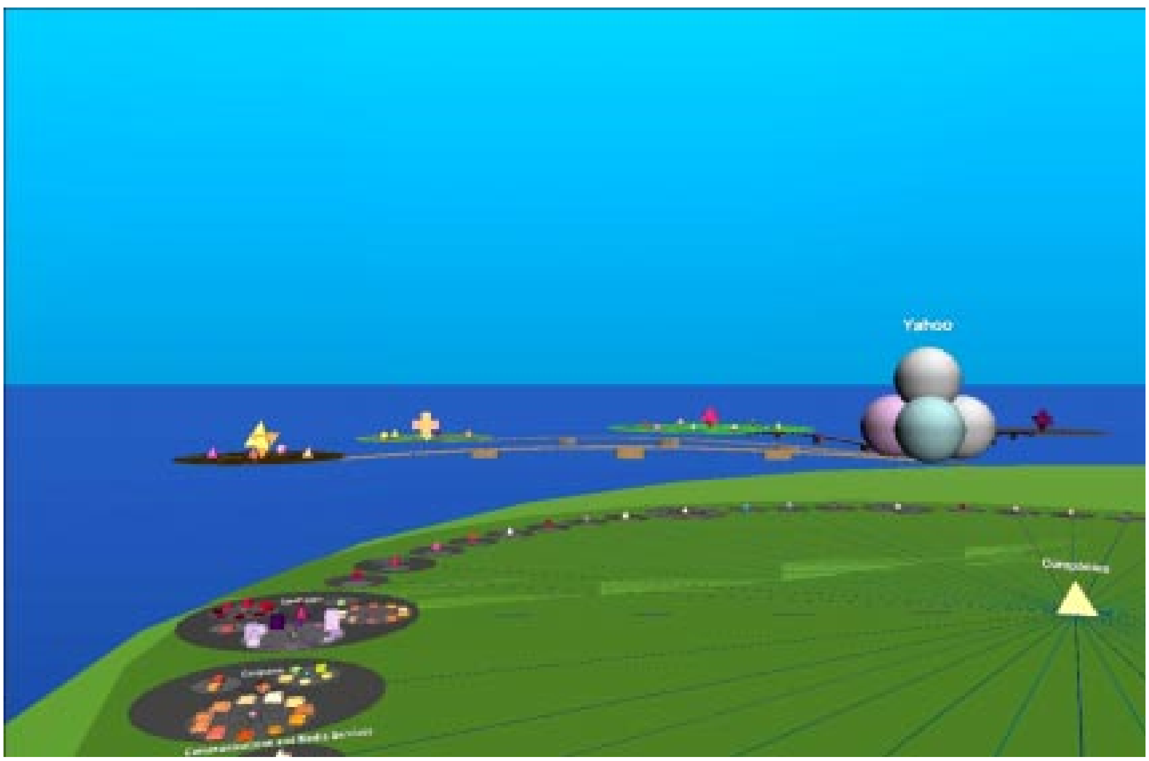

Figure 5.1. A view of the "efficient" Day World.

The goal of Study 3 was to continue the line of research from earlier studies, while focusing specifically on validation of the new visualization model.

After Study 2, many usability and performance problems in the prototype design were fixed, largely based on participants' comments. First, the more structured CityScape layout algorithm was used, to take better advantage of perceptual cues in desktop VR (Keshkin and Vogelmann, 1997). Such cues included radial axes to trace parent-child relationships; circles in the ground plane to delineate category boundaries; variable circle arcs to reflect category sizes; and circle centers with objects and labels to convey category identities. From a design perspective, these visualization elements correspond to Lynch's recommended urban elements. The CityScape algorithm was thus less naturalistic, but more efficient and perhaps more aesthetically satisfying. Second, many usability problems in the prototype were fixed (e.g., by adjusting the brightness and contrast of objects and text, choosing a new "Home" (default) viewpoint for user navigation, and optimizing the VRML scene graph for viewing objects and text at a distance). Third, non-information bearing virtual objects were removed, and the rest geometrically streamlined, to improve graphics performance. (The resulting world was approximately 12% smaller in radius.)

5.2 Methodology

Eight participants were recruited for this study at Umeå University. They were undergraduate and graduate students, primarily in the Departments of Informatics and Computer Science. (The age range of participants was approximately 20-40, and approximately one quarter of them were female. All participants were fluent in English, as required by the Swedish educational system and confirmed in conversation with the experimenter.) The primary criterion for selection was Web browsing experience. This criterion automatically included three factors – computer experience; proficiency with a graphical user interface (windows, icons, mouse, pointer, and scroll bars); and familiarity with large, hierarchical information structure.

The study was a one-factor within-subjects design with two levels of the factor. This factor was virtual world design: "naturalistic" or "efficient" versions of the Day World. Exposure order was fully counterbalanced, with half of the subjects using the naturalistic version first and the other half experiencing the efficient version first. (Details of the experimental design can be found in Appendix F.)

5.2.1 Apparatus

During the study, participants used a standard desktop PC to explore two 3D model worlds. As mentioned, these worlds were naturalistic and efficient versions of the Day World from Study 2. The PC had a Dell 400 MHz Pentium CPU, 128 MB RAM, a Diamond Viper V550 graphics card, a 21" Trinitron monitor, keyboard, speakers, and mouse. The model worlds were viewed in the CosmoPlayer plug-in from SGI for

Web browsers. As in earlier studies in this project, the CosmoPlayer was controlled by moving the mouse on a pad, in conjunction with occasional keyboard modifier keys (e.g., shift, control).

5.2.2 Procedure

Each experimental session lasted approximately 90 minutes per participant, including a rest break. The session began with an introduction to the study, general research context, and lab equipment to be used. This introduction took approximately five minutes.

The second stage of the session was devoted to explanation of experimental procedure, as well as training and practice in the user interface for the study. First, participants received a detailed explanation of experimental procedures, which took approximately five minutes. Next, participants learned the experimental UI. The training included use of the VRML browser's "Home" button, which was always available, in light of research findings that users of large information environments tend to exhibit strong homing behavior (Catledge and Pitkow, 1995; Tauscher and Greenberg, 1997). Participants were instructed not to use the VRML browser's history mechanism (for backtracking), in order to focus their attention more directly on the navigational environment. The first five minutes of this learning time were devoted to basic navigation techniques and practice. Two additional five-minute blocks added more advanced techniques. Such practice involved free exploration of a virtual world that represented sections of the city of Helsinki, Finland. The goal of the introduction, training, and practice was to teach each participant the necessary techniques for performing the experimental tasks successfully, as well as achieving psychological comfort. The participant was advised that questions about experimental procedures, UI techniques, and document language would be permitted and answered, if possible, at any time during the experimental session.

During the third stage of each session, each participant engaged in a scavenger hunt for 20 minutes, and then completed a questionnaire about the visualized data. The investigator began each hunt by locating the participant in the center of the first model world (naturalistic or efficient), and then handing the participant a target description on a paper card. (For the study, 25-target sets had been randomly selected from the visualized data set.) Each card described an informational context (e.g., "Broadway Musicals"), and then presented a specifically-named target (e.g., "Fiddler on the Roof"). Participants were allowed approximately five minutes to find each target, after which they were strongly encouraged to proceed to the next target; participants were requested to find as many targets as possible in the time allowed. When a target was found, the participant was to point it out to the investigator for confirmation. (This confirmation consisted of a spoken word or phrase, such as "Yes," "That's it," or "Correct." Such confirmation was simpler to implement than the musical feedback of Studies 1 and 2. The benefits of this simplification for experimental reliability were thought to outweigh the potential impact of experimenter feedback on participant attitudes. In any case, direct comparison of performance results between studies in this thesis will generally be avoided.) If a target was not found after approximately five minutes, the investigator gave the participant a new target card, and the cycle was repeated. If a target proved difficult, a participant was permitted to skip it and proceed to the next one. Each participant had the option to return to the center of the model world after each target card (indeed at any time), or to remain in their current location. After 20 minutes, the hunt was terminated. The participant was then asked to classify five hypothetical targets on a paper classification form offered for this purpose. In creating this form, items were randomly selected from set of targets not used in the hunt. (A sample form is shown in Appendix F.) The form also asked participants to rate the top-level data categories for perceived size, and to estimate the duration of the scavenger hunt. Participants were allowed as much time as necessary to complete this form. The third stage of the session was followed by an optional five-minute break.

The fourth stage of the session used a different model world (efficient or naturalistic) than the second stage, but was otherwise similar, with the addition of a final 1-page questionnaire. In this questionnaire, participants were asked to subjectively rate the model worlds for apparent size, sense of presence, ease of use, efficiency, and enjoyment. Upon completion of this questionnaire, each participant was paid $15 for participation in the study. The participant was also debriefed by the experimenter about the purpose of the study.

A complete experimental script and related questionnaires can be found in Appendix F.

5.2.3 Measures

The following measures were used in analyzing the results of Study 3.

- **Independent Variables**
  - *BLOCK*: block of trials (1, 2)
  - *DESIGN*: virtual world design (naturalistic, efficient)

- **Dependent Variables**
  - *TARGETS*: number of search targets found, i.e., successful trials (0 - 16)
  - *DURATION*: estimate of block duration (0 – 30 minutes)
  - *SIZE*: estimate of world size (1= small to 5=large)
  - *PRESENCE*: self-reported “sense of presence” (1= low to 5=high)
  - *EASE*: self-reported “ease of use” (1= hard to 5=easy)

– *EFFICIENCY*: self-reported “efficiency” (1=inefficient to 5=efficient)
– *ENJOYMENT*: self-reported “enjoyment” (1= low to 5=high)
– *STRUCTURE*: score on post-block structure test (0 - 1.0)

For scoring a post-block structure test, the primary criterion for correctness was this: accuracy sufficient to guide correct navigational choice in the block’s virtual world. Each structure test contained five items. Each item’s answer was scored for correctness of the first two nodes named. For each correct first node, matching one level below the root of the data hierarchy, a point was added. For each correct second node, matching two levels below the root of the data hierarchy, a point was added. Users could receive a maximum score of two per test item, and thus a maximum total score of ten for all five items. Each test's raw score (0-10) was normalized to a maximum value of 1.0.

For example, one structure-learning test contained the following item: “Folk musician Jeff Buckley.” The correct answer was as follows: “Entertainment → Music → Artists → By Genre → Folk.” One participant offered the following answer: “Entertainment → Music → Folk Music → Artists → Jeff Buckley.” For the node “Entertainment,” this answer received a point; for the node “Music,” this answer received a second point.

5.3 Results

This study sought to identify factors, relationships, and effects that held across the two visualization environments, in order (1) to iterate the designs with regard to usability and (2) to understand the key factors of user response in this research domain. For this reason, study results will generally be analyzed across the experimental conditions, rather than within specific conditions.

5.3.1 Description

	N	Minimum	Maximum	Mean	Std. Deviation
TARGETS	16	2.00	16.00	8.0625	5.2341
DURATION	16	10.00	30.00	18.6250	7.7190
SIZE	16	2.0	5.0	3.250	.753
PRESENCE	16	1	4	2.75	.93
EASE	16	1	5	2.94	1.18
EFFICIENCY	16	1	5	2.94	1.18
ENJOYMENT	16	1.0	4.5	3.406	.987
STRUCTURE TEST	16	.40	.80	.6625	.1310

Table 5.1. Descriptive statistics for the Study 3 measures.

The subjective measures generally clustered about the central value of the five-point scale, with a high of 3.4 for enjoyment. In terms of self-reported measures, participants rated the virtual worlds overall as having a moderate sense of presence. Participants also rated the worlds as having moderate ease of use and efficiency, and moderate to good enjoyment. There was a relatively large number (6) of “Hard” or “Very hard” ratings for the ease of use category.

The distribution for number of targets found was positively skewed, with the majority of observations in the lower levels of performance. The distribution of the post-block structure test was somewhat negatively skewed.

Estimates of block duration were generally slightly less the actual duration of 20 minutes. (There were no significant experimental effects on block duration or world size, so these measures will not be discussed further.)

Detailed descriptive statistics for the experimental measures can be found in Appendix G.

5.3.2 Underlying Factors of Performance

In order to examine the relationships between response measures, a matrix of pairwise correlations was derived (Table 5.2). The four subjective measures (sense of presence, ease of use, rated efficiency, and overall enjoyment) all correlated with each other ($p < .05$), except for sense of presence with ease of use. In addition, ease of use was significantly correlated with number of targets found ($p < .05$). There were no other significant correlations.

	TARGETS	PRESENCE	EASE	EFFICIENCY	ENJOYMENT
TARGETS	1.000				
PRESENCE	.113	1.000			
EASE	**.475**	.288	1.000		
EFFICIENCY	.292	**.530**	**.761**	1.000	
ENJOYMENT	.163	**.517**	**.595**	**.681**	1.000

Table 5.2. Correlations for the Study 3 measures (Pearson, one-tailed, N = 16), bold font shows $p < .05$.

The factor structure underlying the performance-related and four self-reported variables was uncovered with factor analysis. Using the scree plot (Figure 5.2), one factor was selected that accounted for approximately 57% of experimental variance. I will first examine the effect of different experimental conditions on this discovered factor, and then examine the effects for each of the original dependent variables.

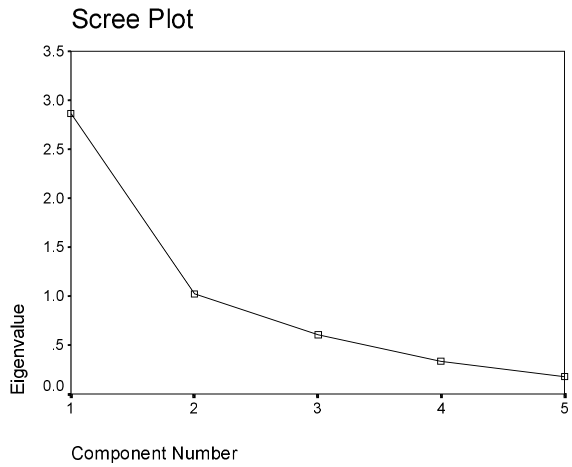

Figure 5.2. Extracted factors vs. variance explained in Study 3.

Varimax rotation was used to enhance the internal probability of the solution, and the resulting factor loadings for the rotated solution are shown in Table 5.3.

	Component
TARGETS	-.01
PRESENCE	**.81**
EASE	**.57**
EFFICIENCY	**.81**
ENJOYMENT	**.85**

Table 5.3. Rotated component matrix for Study 3, one-factor solution (bold font shows factor loadings above .5 or below -.5).

Factor 1 is composed primarily of subjective ("feeling") measures – sense of presence, ease of use, rated efficiency, and overall enjoyment, as indicated by the high loading for these measures varying in absolute value between .57 and .81.

In order to examine the effect of world design on the underlying factor, I performed two one-way ANOVAs using design as a within-subjects measure. There was a significant effect of design on Factor 1 ($F_{1,7} = 6.95$, $p = .034$). Participants felt better after using the "efficient" condition than after the "naturalistic" condition, as shown in Figure 5.3.

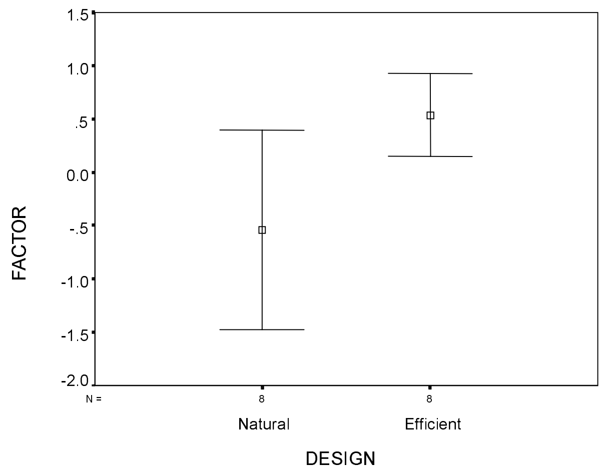

Figure 5.3. Effect of world design on extracted factor 1 in Study 3.

In order to examine the effect of task learning on the underlying factor, I performed a one-way ANOVA using block number as a within-subjects measure. There was no significant effect of task learning on the underlying factor.

5.3.3 Effects on User Strategy

The next analysis examined the particular experimental measures that describe performance and attitude. This inquiry will be presented as effects of three experimental variables: world design, task learning, and (domain) structure learning ability. Statistical analyses will be used to identify significant interactions and main effects involving these variables, as a step towards constructing a model of user behavior and attitude.

5.3.3.1 Effects of World Design

In order to investigate the effect of world design on dependent measures, I performed a series of one-way ANOVAs (one for each of the dependent variables) using world design as a within-subjects measure.

There was a significant effect of world design on ease of use ($F_{1,7}$ = 12.45, p = .01). Participants rated the “efficient” design (3.75) as easier to use than the “naturalistic” design (2.13).

There was also a significant effect of world design on rated efficiency ($F_{1,7}$ = 12.45, p = .01). Participants rated the “efficient” design (3.75) as more efficient than the “naturalistic” design (2.13). (The fact that the F-ratio above is identical to the previous one is a statistical coincidence. The data and analysis were checked thoroughly to ensure that no error had been made.)

There was also a borderline significant effect of world design on overall enjoyment ($F_{1,7} = 3.72$, $p = .095$). Participants rated the "efficient" design (3.88) as more enjoyable than the "naturalistic" design (2.94).

There were no other significant effects of world design.

5.3.3.2 Effects of Task Learning

In order to investigate the effects of task learning on experimental measures, I performed a series of one-way ANOVAs using block number as a within-subjects measure.

There was a significant effect of task learning on number of targets found ($F_{1,7} = 7.0$, $p = .021$). Participants generally found more targets during the second block (10.0) than during the first block (6.13).

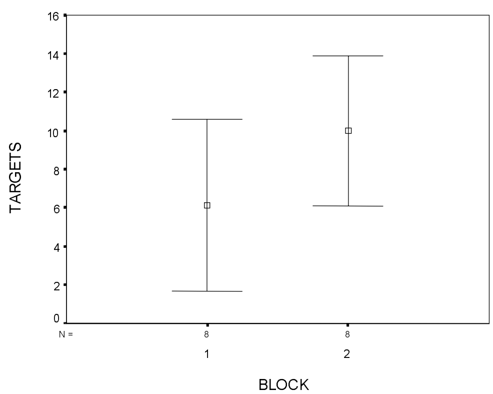

Figure 5.4. Effect of task learning on number of targets found.

There were no other significant effects of task learning.

5.3.3.3 Effects of Structure Learning Ability

In order to investigate the effects of structure learning ability on dependent measures, I averaged each participant's two post-block classification scores, and then aggregated these averages into two groups (high and low). This aggregation created a between-subjects pseudo-factor called STRUCT, which is intended to capture each participant's ability to learn domain information, regardless of effects from virtual world design or task learning. (Speculation suggests that structure-learning ability may be analogous to visual or map memory, but this is a topic for future research.) I then performed a series of one-way ANOVAs using structure learning ability as a between-subjects measure. There were no significant effects of structure learning ability.

5.3.5 Experimental Control

In order to examine effects of asymmetric transfer between experimental conditions, I performed a series of two-way ANOVAs using block number as a within-subjects measure and (1) world design as a within subjects measure or (2) structure learning ability as a between-subjects measure. There were no significant effects of these interactions, and thus no evidence of asymmetric transfer.

To examine any effects possibly hidden by task learning, I restricted the data to the first block only, and performed a series of one-way ANOVAs using world design and structure learning ability as between-subjects measures. There were significant effects of world design on ease of use ($F_{1,6} = 12.0$, $p = .013$), rated efficiency ($F_{1,6} = 24.0$, $p = .003$), and overall enjoyment ($F_{1,6} = 7.71$, $p = .032$). As previously discussed, these effects were all significant or borderline significant for overall experimental sessions as well. The significance of these effects for both within- and between-subjects analysis tends to confirm their presence in the experimental conditions. There were no other significant effects of block-one analysis.

5.4 Discussion

In comparing this study's results with those for hypertext in Study 2, the largest difference was in performance. We can speculatively compare the eight efficient VR blocks in Study 3 with the hypertext blocks in Study 2. These two studies used the same methodology, and exactly eight of the hypertext blocks matched the efficient VR blocks on (1) data set and (2) ordinal position in the session. Since the blocks in Study 2 lasted 15 minutes, only the first 15 minutes of the 20-minute blocks in Study 3 were considered. Analysis showed no significant difference in performance, though the hypertext performance mean (40%) was higher than the VR performance mean (25%).

Observation, however, suggests possible benefits of the efficient VR design. Participants often returned to distant virtual locations by recognizing visual structures, when labels were too distant to read. Such user recognition of visual features suggests VR's power for memorability and wayfinding. Moreover, users expressed pleasant surprise at the new design. It struck several of the participants as a viable representation of Web structure, which offered some of the experiential engagement of games.

5.5 Conclusions

In Study 3, the connection between performance and attitude is ambiguous. Users' performance was not strongly linked to subjective rating, except in the case of ease of use. In Study 3, users seemed to have a relatively objective awareness of each design's suitability for task performance.

Although user performance did not vary between design conditions, the “efficient” design was superior for ease of use and rated efficiency (and perhaps enjoyment). This subjective preference may indicate less cognitive work performed by users of the “efficient” design, relative to the “naturalistic” design. Over a period of extended use, such cognitive improvement would have likely benefits for lack of fatigue and higher motivation. These speculations return to a fundamental goal of HCI design - to reduce the gap between system and user for activities of interpretation and execution (Hutchins et al., 1986).

The apparent improvement in cognitive appropriateness of the “efficient” algorithm calls into question the striving for naturalism in VR design. Study 1 found that users expected a minimum level of naturalism to accommodate the spatial navigational metaphor. Study 3 shows limits on such naturalism for information visualization. The trade-offs between naturalism and abstraction constitute a potential thread of future research in VR information visualization.

Finally, the value of precise “Lynchian” design elements in the superior “efficient” design suggests that wayfinding guidelines from the real world apply to cognitive tasks in electronic worlds, under certain circumstances. Recalling Passini’s (1984) model of iterative wayfinding, the process consists of three repeated stages: mental mapping, plan formulation, and plan execution. The results of Study 3 suggest that large, complex visualization environments share requirements with the physical world for design elements to support this process. That is, visualization environments should provide clear overview for mental mapping; visual orientation and cost-benefit representation to support decision making; and clear affordances for decision execution. Further research is required to identify and define the parameters for successful design of such navigable environments.

Chapter 6

Study 4 - Map-view vs. Fly-through

6.1 Introduction

Study 2 found better search performance with hypertext than with naturalistic desktop VR visualization of hierarchical data. Study 3 demonstrated how the CityScape algorithm could improve the effectiveness of VR visualization. For Study 4, it was decided to investigate the role of the third dimension in the visualization model. Accordingly, a map-view (or birds-eye) version of the "efficient" Day World was created by specialization (sub-classing) in the world-generation software from Study 3. (Figure 6.1)

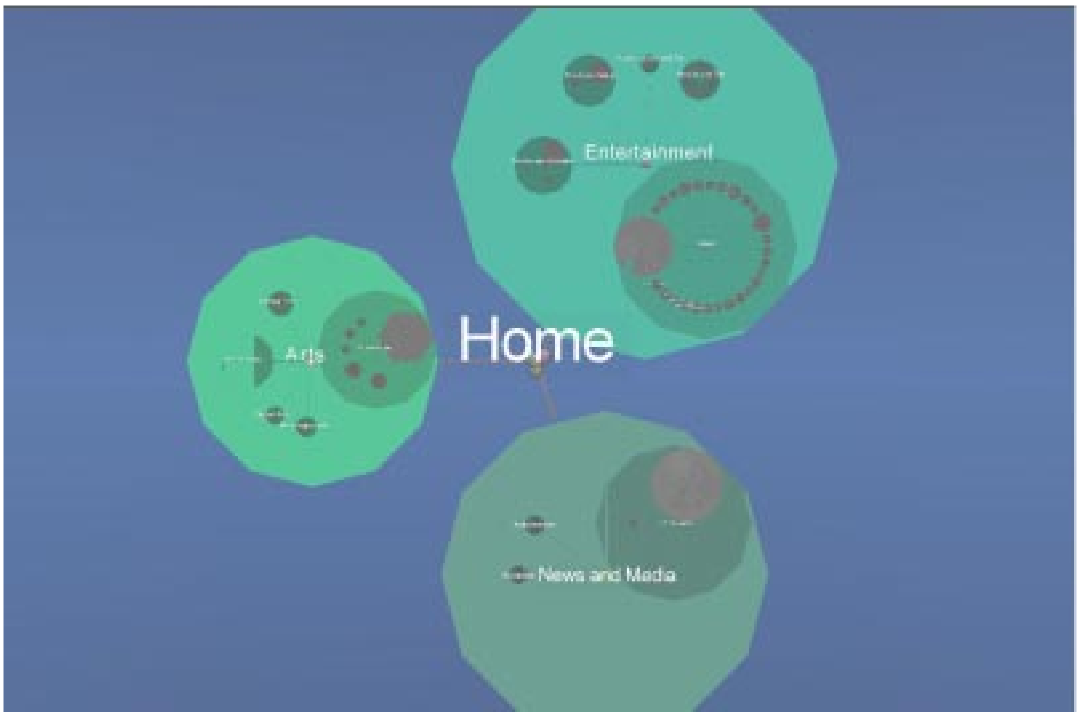

Figure 6.1. A map view of the "efficient" Day World.

The goal of Study 4 was to continue the line of research from earlier studies, while focusing specifically on three issues:

- The effects of fly-through versus map-view presentations of an information landscape,
- Participants' success in learning of information structure
- The effects of individual differences in spatial ability

6.2 Methodology

20 participants were recruited for this study at the University of Toronto. They were upper-division undergraduate, and graduate, students in the Department of Computer Science, the Department of Mechanical and Industrial Engineering, and the Faculty of Information Studies. (The age range of participants was approximately 20-40, and approximately one third of them were female. All participants were fluent in English, as required for enrollment at the University of Toronto and confirmed in conversation with the experimenter.) The primary criterion for selection was Web browsing experience. This criterion automatically included three factors – computer experience; proficiency with a graphical user interface (windows, icons, mouse, pointer, and scroll bars); and familiarity with large, hierarchical information structure. (21 participants actually performed the study, as the seventh participant did not follow the experimenter's instructions and was unable to perform basic tasks in a satisfactory way. The participant met the study's participation criteria fully, however, so the unsatisfactory results could not be anticipated in advance. In any event, it was necessary to discard this participant's results and to recruit another participant.)

The study was a one-factor within-subjects design with two levels of the factor. This factor was virtual world design: map-view or fly-through versions of the "efficient" Day World. Exposure order was fully counterbalanced, with half of the subjects using the map-view version first and the other half experiencing the fly-through version first. (Details of the experimental design can be found in Appendix H.)

6.2.1 Apparatus

Participants used a standard desktop PC to explore two 3D model worlds. These worlds were map and fly-through views of the "efficient" Day World from Study 3. The PC had a Dell 400 MHz Pentium CPU, 128 MB RAM, a Diamond Viper V550 graphics card, a 19" Trinitron monitor, keyboard, speakers, and mouse. The model worlds were viewed in the CosmoPlayer plug-in from SGI for the Web browsers. As in earlier studies in this project, the CosmoPlayer was controlled by moving the mouse on a pad, in conjunction with occasional keyboard modifier keys (e.g., shift, control). A so-called "applet" was written in HTML, Java, and JavaScript to log navigational events in the virtual world to an external file.

In this study, the fly-though design was the same as the "efficient" design in Study 3. The map-view design used the same 3D model as the fly-though design. In the map-view design, however, user navigation was restricted to zooming into and out of the model, as well as 2D translation with reference to the "ground" plane. The map-view thus supported a range of abstraction from detail to overview, and selection of information sub-category at any level.

6.2.2 Procedure

Each participant first completed a standard test of spatial ability. This test was the Minnesota Paper Form Board test. Test administration required approximately 30 minutes, and was carried out sometime (typically a week or so) prior to the experiment.

Each experimental session lasted approximately two hours per participant, including rest breaks. The session began with an introduction to the study, general research context, and lab equipment to be used. This introduction took approximately 5-10 minutes.

The second stage of the session was devoted to (1) training and practice in the user interface for the study, and (2) explanation of experimental procedures in detail. The training included use of the VRML browser's "Home" button, which was always available, in light of research findings that users of large information environments tend to exhibit strong homing behavior (Catledge and Pitkow, 1995; Tauscher and Greenberg, 1997). Participants were instructed not to use the VRML browser's history mechanism (for backtracking), in order to focus their attention more directly on the navigational environment. The first four minutes of this learning time were devoted to basic navigation techniques and practice. Two additional four-minute blocks added more advanced techniques. Such practice involved free exploration of the fly-though virtual world (which was used later in the session for search tasks). The next eight minutes allowed similar training and practice in the map-view version of the virtual world. (The map-view design always followed the fly-through design in the training stage, as this order had proven effective with pilot participants.) The goal of the introduction, training, and practice was to teach each participant the necessary techniques for performing the experimental tasks successfully, as well as achieving psychological comfort. Participants were then informed about the experimental procedures in detail, which took approximately five minutes. The participant was advised that questions about experimental procedures, UI techniques, and document language would be permitted and answered, if possible, at any time during the experimental session.

During the third stage of each session, each participant engaged in a scavenger hunt for 20 minutes, and then completed a questionnaire about the visualized data. The investigator began each hunt by locating the participant in the center of the first model world (map-view or fly-through), and then handing the participant a target description on a paper card. Targets were randomly selected for each participant from the visualized data set, and validated by two engineering students for intelligibility. Target randomization and validation were procedural enhancements based on experience with earlier studies in this project. Each target card described an informational context (e.g., "Broadway Musicals"), and then presented a specifically-named target (e.g., "Fiddler on the Roof"). Each participant was allowed a maximum of five minutes to find each target, and was requested to find as many targets as possible in the time allowed. When a target was found, the participant was to point it out to the investigator for confirmation, and then click an OK button at the bottom of the computer screen. If a target was not found after five minutes, the investigator gave the participant a new target card, and the cycle was repeated. (Unlike Studies 1-3, this study did not allow participants to skip difficult targets. This change was thought (1) to increase experimental consistency and (2) to define an implicit range of search task difficulty to highlight individual differences in performance. Accordingly, direct comparison of performance results between Study 4 and Studies 1-3 should be avoided.) Each participant had the option to return to the center of the model world after each target card (indeed at any time), or to remain in their current location. After 20 minutes, the hunt was terminated. The participant was then asked to classify five hypothetical targets on a paper classification form offered for this purpose. In creating this form, items were randomly selected from the set of those data items not used in the hunt. (A sample form is shown in Appendix H.) The participant was allowed 10 minutes to complete this form. The third stage of the session was followed by a five-minute break.

The fourth stage of the session used a different model world (fly-through or map-view) than the second stage, but was otherwise similar, with the addition of a final 1-page questionnaire. In this questionnaire, participants were asked to subjectively rate the model worlds for sense of presence, ease of use, efficiency, and enjoyment. Upon completion of this questionnaire, each participant was paid $20 for participation in the study. The participant was also debriefed by the experimenter about the purpose of the study.

A complete experimental script and related questionnaires can be found in Appendix H.

6.2.3 Measures

The following measures were used in analyzing the results of Study 4.

- **Independent Variables**
  - *BLOCK*: block of trials (1-3)
  - *DESIGN*: virtual world design (1 = map-view, 2 = fly-through)

- **Dependent Variables**
  - **Original** (from tests, questionnaires, and observation)
    - *MINNESOTA TEST*: number of problems correct on Minnesota Paper Form Board Test (0-64)
    - *TRIALS*: number of search targets attempted[1] (5-29)

[1] In analyzing the navigational logs in Study 4, it proved expedient to measure the number of experimental trials completed, rather than the number of search targets found. Since these two measures had a correlation of 0.994 ($p < .01$, 1-tailed Pearson test, N = 40), it was reasonable to use the trials-completed measure for analysis of the study.

- *STRUCTURE TEST**: score on post-block structure test (0-45)
- *PRESENCE*: self-reported "sense of presence" (1= low to 5=high)
- *EASE*: self-reported "ease of use" (1= hard to 5=easy)
- *EFFICIENCY*: self-reported "efficiency" (1=inefficient to 5=efficient)
- *ENJOYMENT*: self-reported "enjoyment" (1= low to 5=high)

– **Derived**

- From virtual navigation logs
  - *TRAVEL*: distance traveled as percent of virtual world radius, per trial[2] (0 - 100)
  - *PROX*: average zone (circle) of proximity to target (0 = farthest, 4 = closest)
  - *ERRORS*: number of exits from correct zone (circle) of proximity to target, per trial (0+)

- From Minnesota and Structure Tests
  - *SPATIAL*: score on Minnesota Test by quartile (1-4)
  - *STRUCTURE*: average score on Structure Test, by quartile (1-4)

By way of explanation, the study's navigational logs were appended approximately every half-second with the user's current Cartesian coordinates in the virtual world and the time elapsed since the previous log entry. The logs were also appended with an entry when a participant terminated each trial (target search) by pressing an onscreen "OK" button. These logs were filtered by UNIX shell scripts, which also annotated each trial's log segment with layout information from the matching VRML file. Finally, a small C++ utility was written to derive the dependent variables presented above.

For scoring a post-block structure test, as in Studies 2 and 3, the primary criterion for correctness was this: accuracy sufficient to guide correct navigational choice in the block's virtual world. Each structure test contained five items. Unlike in previous studies, all five parts of an item were scored for the correctness of both nodes and links. (This change was made to test a participant's structural learning more completely.) For each correct node name, regardless of position, a point was added (0-5). For each correct link, reflecting two nodes listed in proper order, a point was added (0-4). The total test score (0-45) was the sum of each items' combined node and link scores.

For example, one structure-learning test contained the following item: "History of Hungary." The correct answer was as follows: "Arts → Humanities → History → Regional → Countries." One

[2] An alternative measure was proposed, which was a ratio of the actual distance traveled to the theoretical distance necessary to find all attempted targets, per block. This proposed measure could be called "travel inefficiency." The measure had a mean of 5.14, with a minimum of 1.93 and a maximum of 19.3. Since "travel inefficiency" and distance traveled had a correlation of .963 ($p < .01$, 2-tailed Pearson test, $N = 40$), the simpler distance-traveled measure will be used for analysis of the study.

participant offered the following answer: “Humanities → History → (blank) → Geography → Country.” For the nodes “Humanities,” “History,” and “Country,” this answer received three points; for the link “Humanities → History”, this answer received a fourth point.

6.3 Results

This study sought to identify factors, relationships, and effects that held across the two visualization environments, in order (1) to iterate the designs with regard to usability and (2) to understand the key factors of user response in this research domain. For this reason, study results will generally be analyzed across the experimental conditions, rather than within specific conditions.

6.3.1 Description

	N	Minimum	Maximum	Mean	Std. Deviation
MINNESOTA TEST	40	23	58	43.95	8.48
PRESENCE	40	1	5	3.43	.98
EASE	40	1	5	2.95	.99
EFFICIENCY	40	1	5	3.33	1.05
ENJOYMENT	40	1	5	3.50	1.09
TRIALS	40	5	29	12.27	5.84
TRAVEL	40	.87	6.53	2.31	1.15
PROX	40	1.13	2.46	1.79	.37
ERRORS	40	.45	7.33	2.015	1.52
STRUCTURE TEST*	40	8.00	33.00	22.10	5.84

Table 6.1. Descriptive statistics for the Study 4 measures.

The mean score of approximately 44 for the Minnesota test was slightly below the norm in comparison with a group of 98 engineering students at the University of Michigan in 1970 (who established a mean score of 48.6, with a standard deviation of 6.4). Scores for 998 male engineering freshmen at the University of Minnesota in 1948 were also slightly higher than the present sample (mean 48.7, standard deviation 7.0). In the present study, the minimum score of 23 was at the first percentile relative to the Michigan engineering students, while the maximum score of 58 was at the 95th percentile.

On the Minnesota test, the mean score was 46.00 for males and 39.17 for females. This difference (6.83) is larger than the median (1.18) for eight studies of gender differences in Minnesota test scores (Likert and Quasha, 1970). In the present study, analysis using a one-way ANOVA did not show a significant effect of gender on Minnesota test scores.

The subjective measures generally clustered about the central value of the five-point scale, with a high of 3.5 for enjoyment. In terms of self-reported measures, Participants rated the virtual worlds overall as having a moderate to good sense of presence. Participants also rated the worlds as having moderate ease of use; moderate to good efficiency; and moderate to good enjoyment. There was a relatively large number (15) of “Hard” or “Very hard” ratings for the ease of use category.

The distribution for number of trials completed was positively skewed, with the majority of observations in the lower levels of performance. The distribution of the post-block structure test was negatively skewed, with the majority of observations in the higher levels of performance.

Detailed descriptive statistics for the experimental measures can be found in Appendix I.

6.3.2 Underlying Factors of Performance

In order to examine the relationships between response measures, a matrix of pairwise correlations was derived (Table 6.2). Three of the subjective measures (sense of presence, ease of use, and rated efficiency) all correlated significantly with overall enjoyment ($p < .05$), but generally not with other measures. The objective measures (number of trials completed, travel per trial, errors per trial, average proximity, and measured learning) all correlated significantly with each other ($p < .05$), but generally not with other measures. The subjective measure of efficiency correlated significantly with number of errors ($p < .05$). There were no other significant correlations.

	PRESENCE	EASE	EFFICIENCY	ENJOYMENT	TRIALS	TRAVEL	ERRORS	PROX
PRESENCE	1.000							
EASE	.181	1.000						
EFFICIENCY	.086	.215	1.000					
ENJOYMENT	**.564**	**.407**	**.395**	1.000				
TRIALS	-.057	.011	.174	-.022	1.000			
TRAVEL	.143	-.103	-.232	.046	**-.672**	1.000		
ERRORS	.160	-.223	**-.290**	.026	**-.570**	**.826**	1.000	
PROX	-.094	.079	.258	-.143	**.700**	**-.693**	**-.424**	1.000

Table 6.2. Correlations for the Study 4 measures (Pearson, one-tailed, N = 40), bold font shows $p < .05$.

The factor structure underlying the four performance-related and four self-reported variables was uncovered with factor analysis. Using the scree plot (Figure 6.2), two factors were selected that together accounted for approximately 63% of experimental variance. I will first examine the effect of different experimental conditions on this discovered factor, and then examine the effects for each of the original dependent variables.

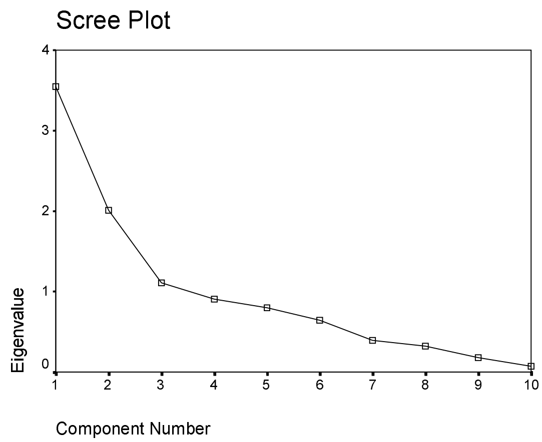

Figure 6.2. Extracted factors vs. variance explained in Study 4.

Varimax rotation was used to enhance the interpretability of the solution, and the resulting factor loadings for the rotated solution are shown in Table 6.3.

	Component	
	1	2
PRESENCE	-.22	**.69**
EASE	.14	**.63**
EFFICIENCY	.35	**.55**
ENJOYMENT	-.10	**.89**
TRIALS	**.83**	.00
TRAVEL	**-.92**	-.01
ERRORS	**-.83**	-.09
PROX	**.82**	-.03

Table 6.3. Rotated component matrix for Study 4, two-factor solution (bold font shows factor loadings above .5 or below -.5).

Factor 1 is composed primarily of objective ("doing") measures – number of trials completed, distance traveled per target, errors per target, and average proximity to target, as indicated by the high loading for these measures varying in absolute value between .82 and .92. By contrast, Factor 2 is composed primarily of subjective ("feeling") measures – sense of presence, ease of use, rated efficiency, and overall enjoyment, as indicated by the corresponding loadings.

In order to examine the effects of spatial ability on the underlying factors, I aggregated participants into quartiles by Minnesota score. This aggregation created a between-subjects pseudo-factor called spatial. I then performed two one-way ANOVAs using spatial ability as a between-subjects measure. There was a significant effect of spatial ability on the "doing" factor above ($F_{3,16} = 4.00$, $p = .027$). Post-hoc analysis using a Tukey test showed a significant difference ($p = .028$) between the first (lowest) quartile and the second quartile, but no significant differences between the second and third quartiles, nor between the third and fourth quartiles. Participants with better spatial abilities generally performed better on the experimental tasks. There was no significant effect of spatial ability on the "feeling" factor.

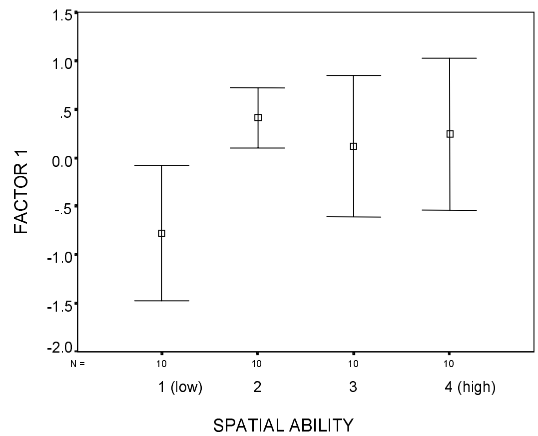

Figure 6.3. Effect of spatial ability on extracted factor 1 in Study 4 (four groups).

On the basis of this result, I focused this analysis on the difference between the lowest and highest spatial-ability groups. There were 10 blocks (5 participants) in the low-ability group and 30 blocks (15 participants) in the high-ability group. There was a significant effect of spatial ability on the "doing" factor ($F_{1,18}$ = 12.32, p = .003). Participants with high spatial abilities generally performed better on the experimental tasks.

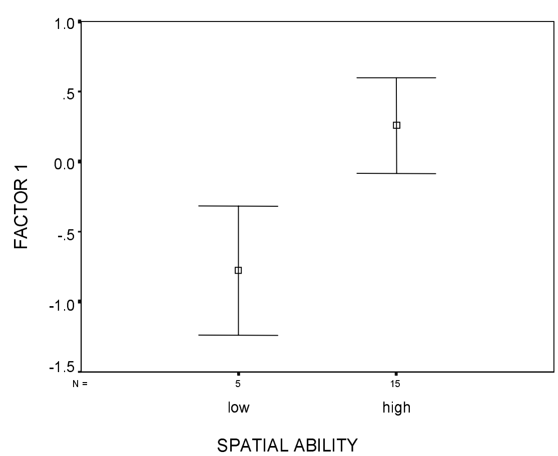

Figure 6.4. Effect of spatial ability on extracted factor 1 in Study 4 (two groups).

In order to examine the effect of world design on the underlying factors, I performed two one-way ANOVAs using design as a within-subjects measure. There was a significant effect of design on the "doing" factor ($F_{1,19}$ = 16.05, p = .001). Participants in the map-view condition performed better on the experimental tasks than in the fly-through condition. There was no significant effect of world design on the "feeling" factor.

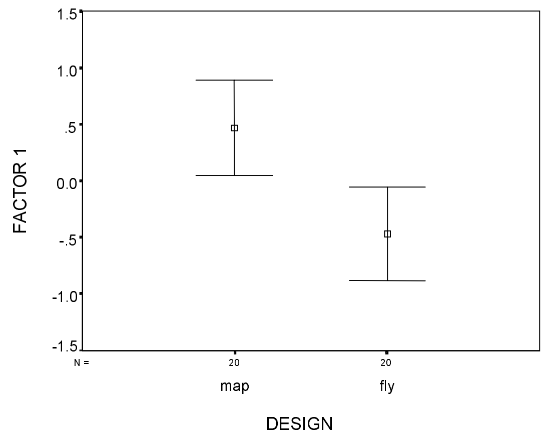

Figure 6.5. Effect of world design on extracted factor 1 in Study 4.

In order to examine the effect of task learning on the underlying factors, I performed a one-way ANOVA using block number as a within-subjects measure. There was no significant effect on either underlying factor.

In general, then, both spatial ability and world design significantly affected underlying performance measures. Underlying attitude, however, was generally unrelated to experimental manipulations or spatial ability, with minor exceptions as noted in later sections.

6.3.3 Effects on User Strategy

The next analysis examined the particular experimental measures that describe performance and attitude. This inquiry will be presented as effects of four experimental variables: spatial ability, world design, task learning, and (domain) structure learning ability. Statistical analyses will be used to identify significant interactions and main effects involving these variables, as a step towards constructing a model of user behavior and attitude.

6.3.3.1 Effects of Spatial Ability

Using the high- and low-ability groups discussed above, I performed a two-way ANOVAs using spatial ability and structure-learning ability as between-subjects measures. There was a significant interaction effect on the number of trials completed ($F_{7,6} = 6.12$, $p < .021$). The performance of users in the lower two quartiles of spatial ability was differentiated less by structure-learning ability than was the performance of users in the higher two quartiles of spatial ability.

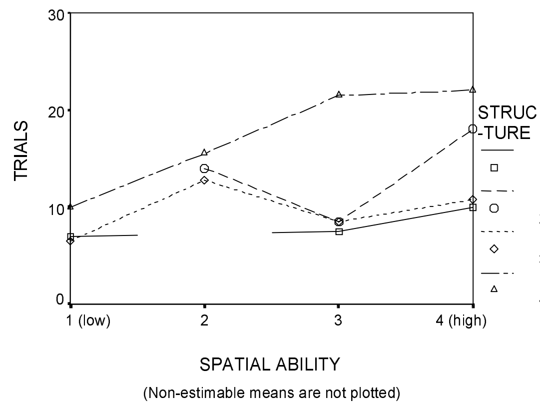

Figure 6.6. Effect of spatial and structure-learning abilities on number of trials completed.

Using the high- and low-ability groups discussed above, I also performed three one-way ANOVAs using spatial ability as a between-subjects measure.

There was a significant effect of spatial ability on number of trials completed ($F_{1,18} = 16.01$, $p = .001$). Participants with high spatial abilities (13.87) completed more trials per participant than people with low spatial abilities (7.5).

There was also a significant effect of spatial ability on virtual distance traveled per trial ($F_{1,18} = 11.18$, $p = .004$). Participants with high spatial abilities (2.01) generally traveled less per trial (more efficiently) than people with low spatial abilities (3.22).

Finally, there was a significant effect of spatial ability on navigational errors per trial ($F_{1,18} = 14.66$, $p = .001$). Participants with high spatial abilities (1.58) generally made fewer navigational errors per trial than did people with low spatial abilities (3.31). There were no other significant effects of spatial ability.

It thus appears that a minimum level of spatial ability was needed to succeed with the experimental tasks, and above this level, other factors presumably account for performance differences.

Note also that two of the three logged behavioral measures –distance traveled and navigational errors – were significantly affected by spatial ability, while proximity to target was not. This difference suggests that non-spatial factors governed proximity to target, as will be discussed in a later section.

6.3.3.2 Effects of World Design

In order to investigate the effect of world design on dependent measures, I performed a series of one-way ANOVAs using world design as a within-subjects measure.

There was a significant effect of world design on number of trials completed ($F_{1,16} = 5.7$, $p = .028$). Participants completed a larger number of trials in the map-view condition (13.6) than in the fly-through condition (11.0).

There was a significant effect of world design on self-reported efficiency ($F_{1,16} = 7.12$, $p = .015$). Participants generally rated the map-view condition (3.75) as more efficient than the fly-through condition (2.9).

There was a significant effect of world design on travel per trial ($F_{1,16} = 7.47$, $p = .013$). Participants generally traveled less per trial in the map-view condition (1.91) than in the fly-through condition (2.7).

There was a significant effect of world design on errors per trial ($F_{1,16} = 6.49$, $p = .02$). Participants generally made fewer navigational errors per trial in the map-view condition 1.5) than in the fly-through condition (2.52).

There was a significant effect of world design on proximity to target ($F_{1,16} = 35.52$, $p = .001$). During each trial, participants were on average closer to each target in the map-view condition (1.99) than in the fly-through condition (1.58).

There were no other significant effects of world design.

6.3.3.3 Effects of Task Learning

In order to investigate the effects of task learning on experimental measures, I performed a series of one-way ANOVAs using block number as a within-subjects measure.

There was a significant effect of task learning on number of trials completed ($F_{1,19} = 8.36$, $p = .009$). Participants generally completed more trials during the second block (13.8) than during the first block (10.8).

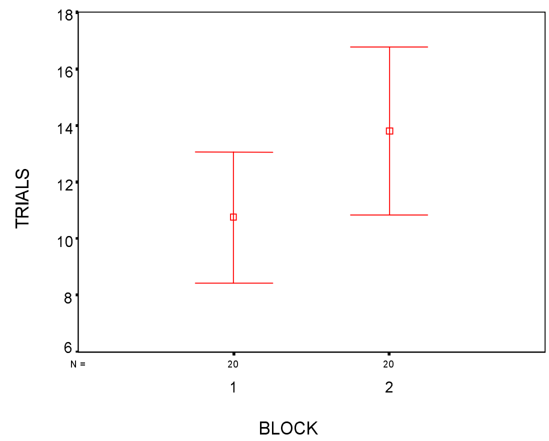

Figure 6.7. Effect of task learning on number of trials completed.

There was a significant effect of task learning on self-reported sense of presence ($F_{1,19}$ = 6.45, p = .02). Participants generally rated the worlds as having a greater sense of presence during the first block (3.85) than during the second block (3.0).

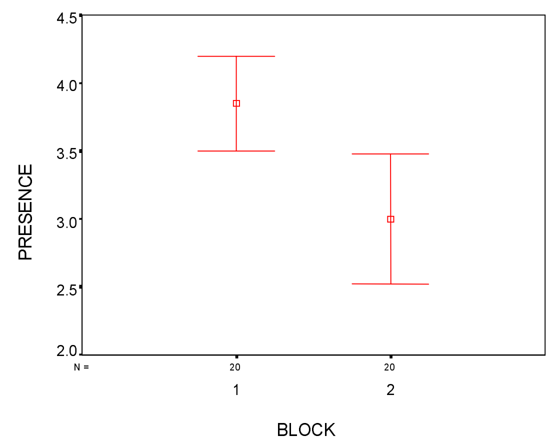

Figure 6.8. Effect of task learning on sense of presence.

There were no other significant effects of task learning.

6.3.3.4 Effects of Structure Learning Ability

In order to investigate the effects of structure learning ability on dependent measures, I averaged each participant's two post-block classification scores, and then aggregated these averages by quartile (e.g., top 25% of average scores in quartile 1). This aggregation created a between-subjects pseudo-factor called STRUCT, which is intended to capture each participant's ability to learn domain information, regardless of effects from virtual world design or task learning. (Speculation suggests that structure-learning ability may be analogous to visual or map memory, but this is a topic for future research.) I then performed a series of one-way ANOVAs using structure-learning ability as a between-subjects measure.

There was a significant effect of structure learning ability on number of trials completed ($F_{3,16} = 7.78$, $p = .002$). Participants with better structure learning ability generally completed more trials.

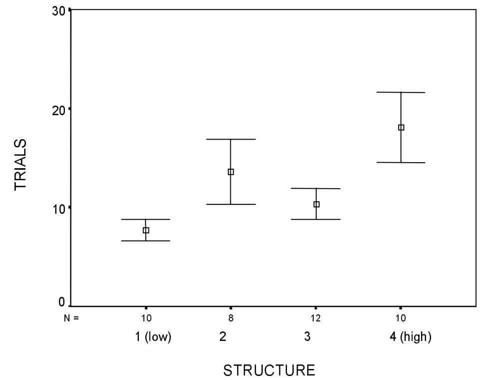

Figure 6.9. Effect of structure-learning ability on number of trials completed.

There was also a significant effect of structure learning ability on proximity to target ($F_{3,16} = 3.98$, $p = .027$). Participants with better structure learning abilities were generally closer to their targets.

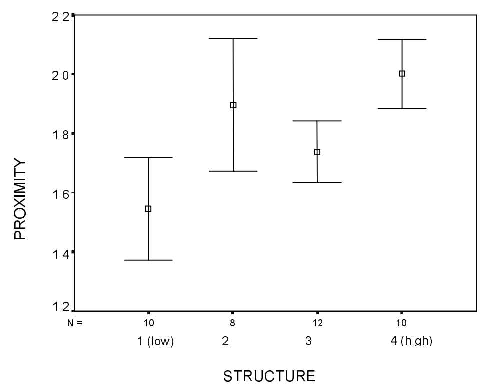

Figure 6.10. Effect of structure-learning ability on proximity to target.

There were no other significant effects of structure learning ability.

As noted earlier, spatial ability did not significantly affect proximity to target, but this section has shown that structure learning ability significantly affected proximity to target. Note that the number of trials completed was significantly affected by both spatial ability and structure learning ability, as well as world design and task learning, which suggests that the number of trials completed depended on a combination of innate user ability, software design, and learning over time. This suggestion will be quantified in a following section.

6.3.3.5 Effects on Proximity to Target

In order to get a clear picture of the time spent by participants at different zones of proximity to targets, I defined five new measures: Z0, Z1, Z2, Z3, and Z4. (Z4 was the zone closest to the target.) Times were expressed as a percentage of each trial time. I then performed a one-way ANOVA using the zones as within-subjects measures.

There was a significant effect of proximity zone on share of trial time ($F_{4,156}$ = 31.04, p = .001). Participants generally spent the most time in the zone farthest from each target, and the least time in the zone second-to-farthest from each target.

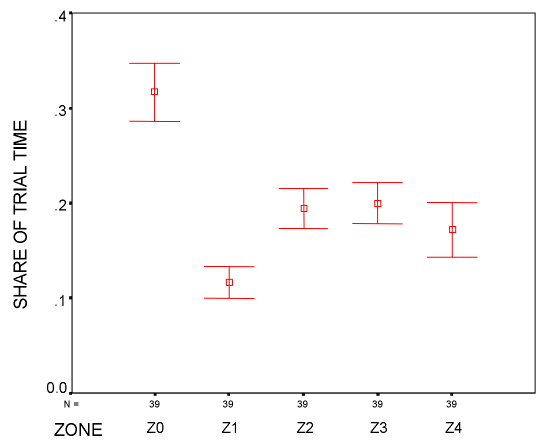

Figure 6.11. Effect of proximity zone on share of trial time.

6.3.4 Learning Effects within Blocks

In order to examine more closely user strategy within each block of trials, I examined the number of trials completed during each two-minute period of a block. There was a significant effect of elapsed time on the number of trials completed ($F_{9,342}$ = 5.82, p = .001).

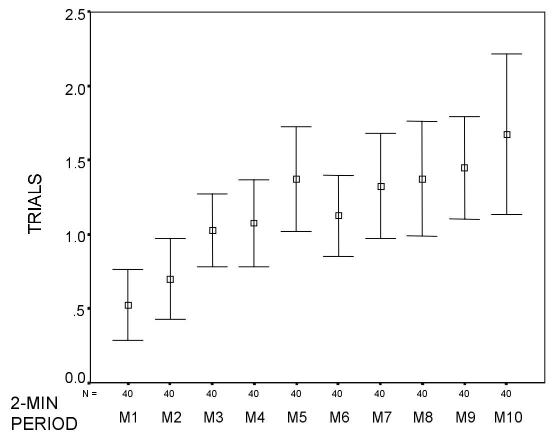

Figure 6.12. Effect of task learning on number of trials completed (per 2-minute period).

6.3.4.1 Effects on Number of Trials Completed

I defined two new measures, H1 and H2. These measures represented the number of trials completed during the first and second half, respectively, of each block. I then performed a series of two-way and one-way ANOVAs similar to the ones discussed above. There was a significant effect of within-block task learning on number of trials completed ($F_{1,39}$ = 41.02, p = .001). Participants generally completed more trials during the second half-block (7.58) than during the first half-block (4.7).

6.3.4.2 Effects in First Half-Block

There was no significant interaction effect of within-block task learning and elapsed time on number of trials completed.

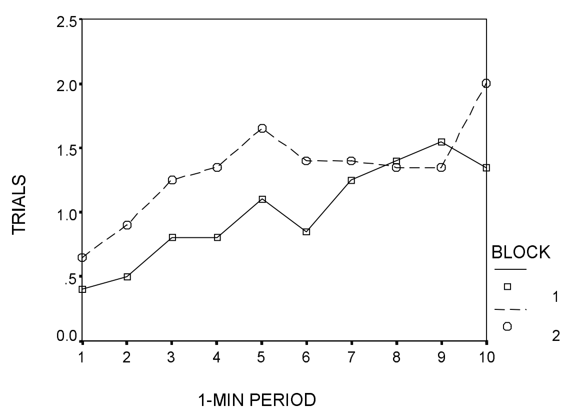

Figure 6.13. Effect of task learning on number of trials completed
(per 1-minute period) in first half-block.

There was, however, a significant effect of inter-block task learning on the number of trials completed ($F_{1,19}$ = 13.25, p = .002). Participants generally completed more trials during the second block's first half (5.8) than during the first block's first half (3.6).

There was no significant effect of world design on trials completed during the first half-block.

6.3.4.3 Effects in Second Half-Block

There was no significant interaction effect of world design and elapsed time on number of trials completed.

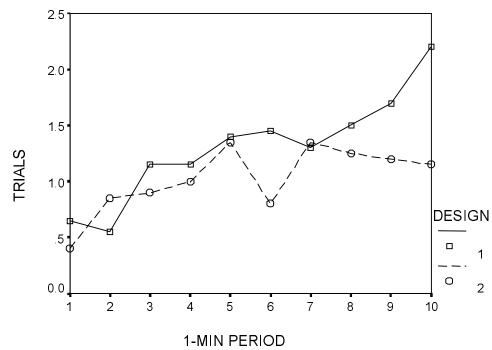

Figure 6.14. Effect of task learning on number of trials completed (per 1-minute period) in second half-block.

There was, however, a significant effect of world design on number of trials completed ($F_{1,19}$ = 10.36, p = .005). Participants generally completed more trials in the map-view condition (8.7) than in the fly-through condition (6.45).

There was no significant effect of inter-block task learning on number of trials completed in the second half-block.

6.3.5 Data Classification

In order to determine which measures could most concisely account for differences between various combinations of settings of independent variables in the experiment, I ran a series of discriminant analyses on the dependent measures discussed above.

For the differences between the two levels of world design, three key variables were identified that together represented 100.0% of experimental variance: proximity to target, self-reported efficiency, and self-reported enjoyment. The corresponding discriminant function coefficients are shown in Table 6.4.

	Function
	1
EFFICIENCY	.819
ENJOYMENT	-.688
PROXIMITY	.652

Table 6.4. Standardized canonical discriminant function coefficients (world design).

The map view was characterized by greater self-reported efficiency and proximity to target, but lower sense of enjoyment (in the first block of trials).

For distinguishing levels of spatial ability, two key variables were identified that together represented 88.8% of experimental variance.

	Function	
	1	2
EFFICIENCY	.827	.758
ERRORS	1.050	-.394

Table 6.5. Standardized canonical discriminant function coefficients (spatial ability).

Participants with lower spatial ability reported higher task efficiency, but actually made more navigational errors.

6.3.6 Linear Regression

In order to explain variance in the number of trials completed on the basis of the measures for spatial and structure-learning ability, I ran a series of linear regressions.

Using all cases, multiple linear regression was carried out with number of trials completed as the criterion variable, and structure learning and spatial ability as the predictor variables. A significant linear relationship was found ($F_{2,37} = 11.31$, $p = .001$) with 38% of the variance accounted for ($R = .62$). The best fitting equation was: trials = 1.99 + 1.56*spatial + 2.5*structural. The corresponding standardized beta coefficients were .30 and .49.

Using just the map-view cases, multiple linear regression was carried out with number of trials completed as the criterion variable, and structure learning and spatial ability as the predictor variables. A significant linear relationship was found ($F_{2,17} = 9.21$, $p = .002$) with 52% of the variance accounted for ($R = .72$). The best fitting equation was: trials = .18 + 2.25*spatial + 3.06*structural. The corresponding standardized beta coefficients were .40 and .54.

In the map-view regression model, it was possible that the lowest spatial-ability group accounted for much of that variance explained by spatial ability. Accordingly, using just the higher-ability group of 15 subjects (as previously discussed), multiple linear regression was carried out with number of trials completed as the criterion variable, and structure learning ability as the predictor variable. A significant linear relationship was found ($F_{1,28}$ = 11.76, p = .002) with 30% of the variance accounted for (R = .54). The best fitting equation was: trials = .5.36 + 3.11*structural. The corresponding standardized beta coefficient was .54.

Using just the fly-through cases, multiple linear regression was carried out with number of trials completed as the criterion variable, and structure learning and spatial ability as the predictor variables. A borderline significant linear relationship was found ($F_{2,17}$ = 3.43, p = .056).

In summary, 38% of the between-subjects variation in task performance was accounted for by differences in spatial and structure-learning ability. Separate analyses of the map-view versus fly-through conditions showed the following: the predictive effect of the two types of ability was much greater for the map condition, where the abilities accounted for over half the variation in performance between individuals. For the fly-through condition, the linear relationship between the two types of ability and performance was not significant at the .05 level. These findings demonstrate how individual differences in types of ability can vary in importance, depending on the style of user interface. In future research, testing UI competence independently of spatial and structure-learning ability might help to account for some of the unexplained experimental variance (Waller, 1999).

6.3.7 Experimental Control

In order to examine effects of asymmetric transfer between experimental conditions, I performed a series of two-way ANOVAs using block number as a within-subjects measure and (1) world design as a within-subjects measure, (2) spatial ability as a between-subjects measure, or (3) structure learning ability as a between-subjects measure.

There was a significant effect of spatial ability and task learning on self-reported ease of use ($F_{3,16}$ = 3.65, p = .035). Participants in the lowest ability group rated the second block as easier to use than the first block. In the (3) higher ability groups, participants' ease-of-use ratings tended to decrease with ability in the first block, and to increase with ability in the second block.

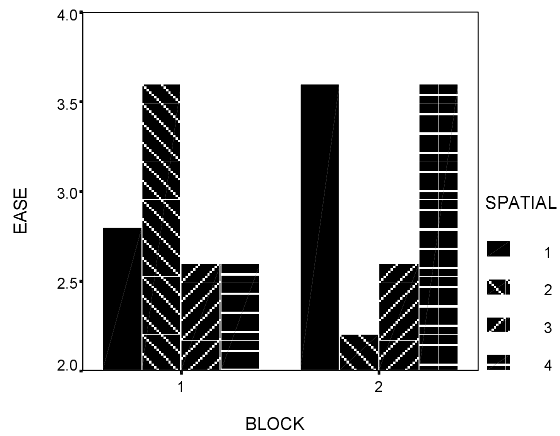

Figure 6.15. Effect of task learning and spatial ability on ease of use.

There was a significant effect of structure learning ability and task learning on trials completed ($F_{3,16}$ = 6.12, p = .006), proximity to target ($F_{3,16}$ = 5.73, p = .007), and rated efficiency ($F_{3,16}$ = 8.98, p = 001). Higher structure-learning ability tended to lead to improved performance.

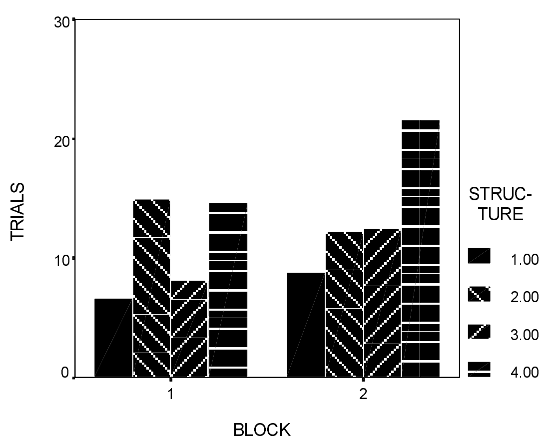

Figure 6.16(a). Effect of task learning and structure-learning ability on number of trials completed.

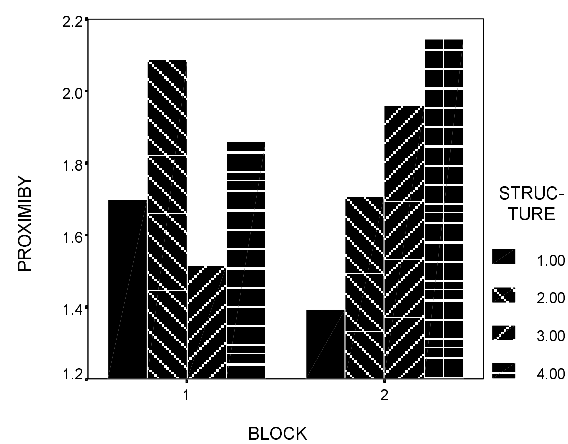

Figure 6.16(b). Effect of task learning and structure-learning ability on proximity to target.

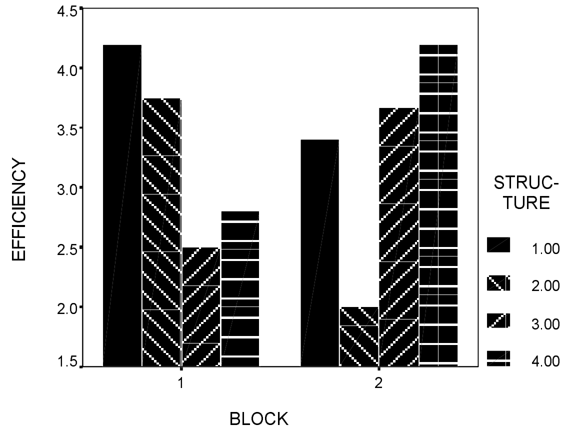

Figure 6.16(c). Effect of task learning and structure-learning ability on self-reported efficiency.

There were no other significant effects of these interactions.

To examine any effects possibly hidden by task learning, I restricted the data to the first block only, and performed a series of one-way ANOVAs using world design, spatial ability, and structure learning ability as between-subjects measures. There was a significant effect of world design on self-reported sense of enjoyment ($F_{1,18}$ = 12.79, p = .002). Participants generally rated the fly-through condition as more enjoyable (4.2) than the map-view condition (3.3).

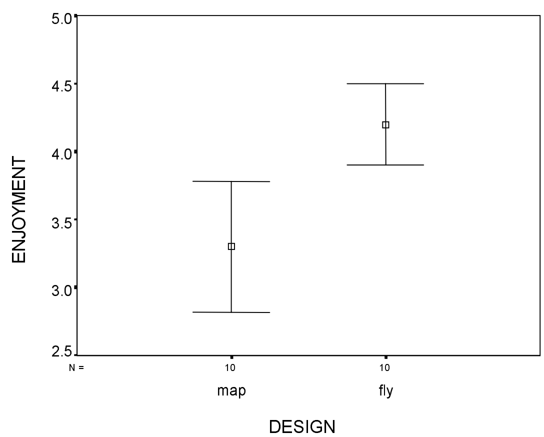

Figure 6.17. Effect of world design on enjoyment (during first block).

There was also a significant effect of world design on proximity to target ($F_{1,18}$ = 12.38, p = .002) in block-one analysis, similar to the effect in all-blocks analysis, but no further effects of world design in block-one analysis.

There was a significant effect of structure learning ability on self-reported efficiency ($F_{3,16}$ = 5.36, p = .01). Participants with higher ability generally rated the virtual worlds as less efficient than did participants with lower ability. This result offers another example of performance not affecting subjective response, since it was previously noted that participants with higher structure learning ability completed significantly more trials over two blocks than did participants with lower ability.

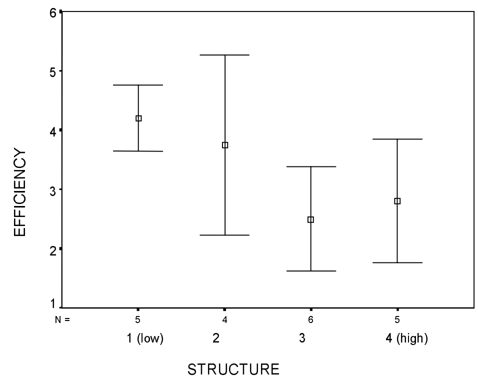

Figure 6.18. Effect of structure-learning ability on self-reported efficiency (during first block).

There were no other significant effects of block-one analysis.

To verify that the spatial and structure learning ability measures represented different skill sets, I analyzed the participants' raw scores on the Minnesota Paper Form Board test and the post-block classification forms using a one-tailed Pearson correlation test. The correlation was not significant.

6.3.8 Overview of Results

An overview of significant results is presented in Table 6.6. These results generally fall into four groups, as discussed below.

First, there were significant effects of world design on five dependent variables – rated efficiency, trials completed, travel, navigational errors, and proximity to target. These effects were captured in the significant effect of world design on the underlying "doing" factor of performance.

Second, there were significant effects of four independent variables – world design, task learning (block), structure-learning ability, and spatial ability - on the number of trials completed. These results highlight the multi-dimensional nature of task performance in this research domain.

Third, there were significant effects of spatial ability on three dependent variables – number of trials completed, travel, and navigational errors. These effects were captured in the significant effect of spatial ability on the underlying "doing" factor of performance.

Measure \ Factor	Design	Block	Struct	Spatial
Presence		*		
Ease				
Efficiency	*			
Enjoyment				
Trials	*	**	**	**
Travel	*			*
Errors	*			*
Proximity	**		**	
"Feeling"				
"Doing"	**			*

Key
** = p < .01
* = p < .05

Table 6.6. Overview of significant results.

Finally, there were significant effects of structure-learning and spatial ability on the log-based factors of analysis. There was a significant effect of spatial ability on travel and navigational errors, while there was a significant effect of structure-learning ability on proximity to target. The significant effects of these abilities were mutually exclusive; this observation suggests that the measures reflected distinct cognitive attributes of the experimental participants.

6.4 Conclusions

The most striking result of Study 4 is the general superiority of the map-view design. The design apparently uses the third spatial dimension more effectively than does the fly-through design, leading to better performance and perhaps greater cognitive efficiency. In terms of Passini's (1984) wayfinding model, the superiority of the map-view design for mental mapping should be clear.

Although Study 4 showed no effect of spatial ability on the subjective measures, this ability had a role to play in user success with search tasks. Closer analysis showed that only users with below-average spatial ability had trouble with experimental tasks, while users with average and above-average ability performed relatively well. This finding echoes that of Campagnoni and Ehrlich (1989), who found that users with relatively weak visualization skills performed poorly on navigational tasks in a hypermedia system. In the present study, users with average and above-average spatial ability were differentiated in performance primarily by structure-learning ability. It may be that such users could achieve optimal performance with future visualization environments via features designed to facilitate structure learning.

As in earlier stages of this research, the connection between performance and attitude is ambiguous. In Study 4, users' ratings of the virtual worlds were not strongly linked to design, except in the case of self-reported efficiency. In this case, users seemed to have a relatively objective awareness of each design's utility for task performance. Otherwise, users did not have a strong subjective preference for either world design, which suggests that it is possible to design worlds that are both efficient and enjoyable to use.

Chapter 7

Conclusions

7.1 Summary

After reviewing related research, this thesis presented four user studies that took place over a period of two years. The studies pursued themes related to the central question of when and how desktop VR might be an effective technique for visualizing hierarchical data.

The first study compared three different versions of a virtual landscape: the Day World, the Dusk World, and the Night World. Subjects generally preferred the Day World version for sense of presence, ease of use, and overall enjoyment. Task performance was unaffected by experimental conditions.

In order to establish a baseline for performance and attitude, the second study compared the Day World design with a hypertext based on the same data set. Hypertext generated better task search performance, as well as higher ratings for ease of use and efficiency, but the VR design was rated as more enjoyable overall.

For the third study, in order to incorporate some of the structural advantages of hypertext into the VR prototype, a more structured layout algorithm (CityScape) was implemented for the Day World. A user study showed the new design to be considered easier to use and more efficient than the old one. Task performance, however, was unaffected by world design in this study.

Finally, to investigate the value of the vertical (third or Z-) dimension of the 3D model, a version of the Day World was implemented that supported navigation on a zoomable, map view of the landscape. This map version generated better task search performance in experimental participants, as well as better ratings for efficiency. Participants in the bottom 25% of standardized spatial ability had significantly worse performance than those in the top 75% of ability (for both fly-through and map-view conditions). Participants with higher structure-learning ability performed significantly better than those with lower ability. Moreover, participant performance improved with task learning. Performance thus reflected four significant factors: spatial ability, structure-learning ability, world design, and task learning.

7.2 Methodology and Spatial/Semantic Issues

In presenting the results of the above user studies, the thesis has attempted to balance technical and psychological perspectives. A task focus was chosen for this purpose. Accordingly, each study presented results in terms of both UI design condition and user response, as mediated by experimental search tasks.

In order to investigate information visualization and user navigation in desktop VR, it was necessary to adapt existing methodologies and/or develop new ones. Five methods were used in this series of studies: search task scoring, attitude questionnaires, spatial ability tests, structure learning questionnaires (classification forms), and VR navigation logging. Search tasks were administered because they seem to occur often in the real world, and because they involve user abilities, task learning, and reportable experiences. The ability to vary data sets and UI conditions independently also made search tasks valuable. Attitude questionnaires are common in HCI research, and they proved an effective vehicle for participants to express estimated duration, sense of presence, ease of use, efficiency, and overall enjoyment. For spatial ability evaluation, the Minnesota Paper Board Form test provided an instrument that revealed significant effects on task performance. Structure-learning ability, as assessed using the classification forms, had significant effects on task performance. Future work is required to characterize more precisely the ability tested by these classification forms. Finally, the use of navigational logs for analyzing user navigational behavior proved valuable in Study 4. Several measures were extracted from the logs: virtual distance traveled per trial, navigational errors, and average proximity to target. These measures correlated closely with number of trials completed, while providing insight into user behavior and strategy.

As mentioned in the introduction to the thesis, the dichotomy between spatial and semantic structure provided the motivation for this research. During the first three studies, this dichotomy was regarded as primarily one of technology, that is, an issue for information design. The first study, on the one hand, showed the importance of adequate spatial cues for a user's sense of presence, ease of use, and rated efficiency. The second study, on the other hand, showed the importance of clearly presented semantic structure, in the form of hypertext, for task performance. The third study, by implementing the CityScape algorithm, resolved these issues somewhat by developing a prototype that users found to be both usable and efficient. Having resolved some of the technical issues, in the fourth study, spatial and semantic issues were tested as user attributes, in the form of spatial ability and structure-learning ability. Results showed a wide variation in user ability and accompanying experimental task performance. In particular, people with the lowest spatial ability had significant difficulties with the user interface; users with higher structure-learning ability had better task performance. With regard to spatial and semantic ability, these results suggest that a single software tool may be usable by a majority of users, but that a minority may be served best by specialized tools.

7.3 Performance, Ease of Use and Efficiency

Factor analyses of the four studies indicated the presence of a “feeling” (ease of use, efficiency, and enjoyment) and/or “doing” (task performance) factor.

During the first two studies, there were large perceptual differences between the UI conditions. In the first study, for example, there were strong differences in the use of lighting and color. In the second study, hypertext and desktop VR differed in dimensionality and the presence of spatial cues. In these studies, “feeling” and “doing” appeared to have been correlated, with better sense of presence (etc.) associated with better performance. For this reason, “feeling” and “doing” were combined in a significant single factor for the first two studies. In the third and fourth studies, the UI contrasts were not as great, as all of the experimental conditions used desktop VR with similar spatial cues, lighting, and colors. Thus it seems reasonable that in the factor analyses carried out on the results of these studies, the earlier “feeling/doing” factor separated into two distinct factors, one for “feeling” and one for “doing.”

These speculations return to a question posed in the introduction to the thesis: namely, to which experimental variables/conditions do users respond objectively or subjectively? In Studies 1 and 3, design conditions were similar enough not to affect task performance. In Studies 2 and 4, design conditions were different enough to affect task performance: hypertext was better than the “naturalistic” Day World”, and the map-view of the “efficient” Day World was better than the fly-through. Both the hypertext and map-view designs featured data structure alignment with the 2D user interface. In these designs, the theoretical gulf between information evaluation and command execution (Hutchins et al., 1986), or the effort in proceeding from mental maps to navigational plan execution (Passini, 1984), may have been significantly reduced. (The 3D user interfaces created greater interaction complexity, which might not have been compensated for by better visual apprehension of information and structure.)

User ratings for ease of use clearly favored one design condition in each of the first three studies (Day World, VR, and “efficient” Day World). By this measure, the map-view and the immersive-view in Study 4 were not significantly different for subjective ease of use. As ratings of this measure did not follow objective behavior in Studies 1, 3, and 4, the measure may have captured an aspect of cognitive work instead, perhaps related to clarity of visual structure perception.

In each study where they were requested, user ratings for the efficiency of design conditions matched objective behavior (hypertext, map-view), rated ease of use (hypertext, “efficient” Day World), or both. Of the subjective measures, rated efficiency seems to reflect the most accurate user awareness of objective behavior. It may thus have captured a user's sense of the difficulty in translating visual perception of structure ("ease of use") into task-based commands for execution ("task performance").

7.4 Sense of Presence, Enjoyment, and 3D

Only in the first study did sense of presence significantly differentiate design conditions, favoring the Day World. In addition, in the fourth study, sense of presence was rated higher in the first block of trials than in the second. It appears that rated sense of presence may have responded to a perceptual factor, either strong lighting and bright colors (as in the Day World), or novelty of the medium when 3D models were similar (as in Study 4).

During each of the first two studies, users clearly favored one design condition for enjoyment (Day World, VR). Yet in the last two studies, a clear preference emerged only in the first block of the fourth study (immersive-view). It may thus be that rated enjoyment was a response to a perceptual factor, perhaps vividness of 3D simulation. In this case enjoyment stands squarely in opposition to the superiority for task performance of hypertext in Study 2 and the map-view in Study 4. There may, in fact, be a design trade-off between the perceptual vividness of games and the task utility of a 2D textual interface. However, the potential motivational benefit of immersion deserves further research. From the point of view of this discussion, an ideal interface might meld the two design conditions by offering both map support and immersive browsing. While this solution has been tried in some computer games, scientific research is needed to investigate this phenomenon further.

As previously mentioned, the most successful UIs of this thesis research were those that aligned the data structure with the 2D user interface (hypertext and the immersive-view). Yet the results of Study 2 showed the "naturalistic" version of the Day World to be more enjoyable than hypertext, so the same result might hypothetically be obtained for later versions of the Day World. Comparison of the results of Study 3 with comparable cases from Study 2 showed no significant difference in task performance. In addition, Study 4 showed that the map-view was better suited for search task performance than the immersive-view. Putting these three results together suggests that the map-view of Study 4 might offer the user benefits for both performance and enjoyment.

7.5 Contributions

Most generally, the answer to the original research question is as follows: it appears that the map-view of the "efficient" version of the Day World design is an effective user interface for the top 75% of users in spatial ability for performing simple searching tasks in hierarchical information visualization in desktop VR. In reaching this answer, the current research has made three general contributions in the areas of user behavioral modeling, system design, and usability evaluation:

1. Identified and partly quantified psychological aspects of users that affect search task performance and resulting attitudes:
    a. Determined that a minimum of spatial ability (25th percentile on Minnesota Paper Form Board Test) is required for successful completion of search tasks in desktop VR
    b. Found that higher structure-learning ability significantly improved search task performance in hierarchical information visualization
    c. Identified significantly different factors of objective and subjective user response to experimental conditions

2. Iteratively designed, implemented, and evaluated a software prototype for usability and robustness:
    a. Developed a usable prototype of desktop VR visualization for general hierarchical information, and empirically verified prototype's usability for simple search tasks
    b. Implemented the "CityScape" layout algorithm for structural clarity, in a visualization prototype with approximately 1500 information nodes on six levels of hierarchy
    c. Found that environments that place data structure in correspondence with a 2D UI generated significantly better search task performance than environments that place data structure in a desktop 3D model.

3. Developed and refined a research methodology for investigating tasks, users, and software in desktop 3D environments:
    a. Developed test instrument for structure-learning ability, and applied this test, and a standardized test of spatial ability, in interpreting experimental outcomes
    b. Developed and used tool for logging of navigational activities in virtual-world coordinate system, and developed analytical constructs to interpret these logs in terms of experimental issues
    c. Developed structured search tasks based on a standardized data hierarchy that allowed users with significant individual differences to use common apparatus, procedures, and data collection methods.

7.6 Future Work

Future work proposed for this area can be organized under the three categories of research contributions above.

7.6.1 User Performance and Attitude

While the current research began to explore psychological interactions between user responses under controlled circumstances, further research is required to characterize these interactions more deeply. Characterization would be particularly useful for interactions between spatial and structure-learning abilities, as well as trade-offs between perception and cognition in user awareness. Such interactions could also benefit from investigation in a more diverse participant population and/or a longitudinal study.

7.6.2 Visualization Prototypes

Future work is also necessary to clarify the design parameters discussed above for information visualization, as well as to investigate potential integration of the experiential benefits of immersive views with the effectiveness benefits of overview. In particular, it would be worth investigating other potential algorithms for structural layout. Furthermore, issues of software artifact scalability, evolution, and distribution need to be addressed, in order for the present prototypes to be developed into usable tools. The World Wide Web is a clear candidate application domain; perhaps even more so are corporate and academic "intranets", which can be designed and managed in an integrated fashion by IT personnel. Other possible lines of inquiry concern the suitability of landscape visualizations of hierarchical data for dynamic, on-the-fly information, as well as the personalization of these visualizations for users of different abilities, personalities, and other attributes. In 3D visualization environments, there is a potential trade-off between (1) stable, familiar information structure (potentially long-term) that users must recognize, and volatile, unfamiliar elements (potentially short-term) whose changes must be emphasized; a topic for future research is the design of techniques to handle such trade-offs appropriately.

Further considerations arise for extending the application functionality of the research prototype. From a temporal point of view, the tool requires support for both retrospective and forward-looking navigation. That is, the tool would require (1) some sort of bookmark or "breadcrumb" mechanism to support returning to a previously-visited virtual location and (2) support for guidance towards a location of interest, perhaps in the form of a highlighted virtual route or metaphorical rapid transportation. (The choice of an appropriate metaphor for such transitions is a research topic in its own right.) Related to the notion of navigational support are appropriate tools to assist users with understanding the virtual environment itself, e.g., maps and queries. Maps, first, are a standard feature of 3D desktop computer games; choosing an appropriate strategy for integrating a map with information visualization (e.g., inset, popup, or paned map) would be an interesting research topic. Especially in virtual reality, it may be possible to further integrate detail and overview through semantic and/or physical zooming. A query mechanism, second, would be useful to augment the visualization's browsing features with rapid access to specific topics of interest. Presuming a search engine of appropriate architecture and functionality, the research task would be to effectively integrate the query mechanism, particularly output, with the existing visualization. Dynamically emphasizing aspects of the visualization landscape to reflect query results is one promising line of inquiry (e.g., Chen, 2000; Rojdestvenski, Modjeska, Pettersson, Rojdestvenskaia, & Gustafsson, 2000).

7.6.3 Research Methodology

With regard to the current thesis' methodology, further research might be desirable for several aspects. First, research on information visualization and user navigation needs to be generalized beyond search tasks, to consider browsing and other tasks. Second, ambiguous or hard-to-interpret questionnaire items need to be revised and re-evaluated. Third, a larger participant pool would be needed to validate the appropriateness of standardized tests in the VR domain; in addition, the classification forms to assess structure-learning ability need to be refined for clarity and time efficiency, and the underlying ability tested by this instrument needs to be characterized more precisely (as previously noted). Finally, the navigational logging and analysis could benefit from a larger set of parameters (e.g., UI events) and additional analytical constructs (e.g., user orientation relative to target).

Beyond the previously-discussed research questions, more general ones arise. First, there are trade-offs between handcrafted and algorithmic solutions to information representation. For example, to what extent would 2- or 3D icons improve user navigation and learning, and how can such icons be efficiently incorporated into VR visualization? Second, given the novelty of the medium, another potential line of inquiry concerns the immersion offered by different hardware and software. What effect would such equipment have on the issues in the thesis? Finally, it is worth considering visualization of very large-scale, dynamic structures such as the Web, and for social navigation of shared 3D visualizations by physically distant collaborators.

7.7 Final Comments

Identification of key psychological attributes has implications for research on computer users, in that this work shows several psychological responses interacting under controlled circumstances. The interaction of spatial ability and structure-learning ability is particularly interesting in this regard, as is the interaction between objective and subjective responses.

The thesis research has extended existing human-computer interaction methodologies into a relatively new domain, in which there is a small but growing amount of empirical work. In the thesis studies, subjective questionnaires were found useful. For this research, a search task was developed from a large, standardized information source – the Yahoo! Web index – which suggests the possibility of adapting other information sources for related research. Two user ability tests were used, a standard one for spatial ability and a novel one for structure-learning ability, in ways that may be useful for research on searching behavior in other types of information environment. For VR research in particular, this project's exploratory method for navigation logging and analysis could be extended and refined for other research contexts, e.g., CSCW work on virtual communities, military work on combat simulations, and perhaps even urban design work on pedestrian movement.

Finally, the software prototypes developed in this thesis research have implications for information technology, particularly in the area of visualization. In Study 2, comparison with hypertext emphasized the need for clear visual structure for users' evaluation and execution. Studies 1 and 3 underscored the importance of good visual design for user satisfaction (sense of presence, ease of use, and enjoyment) and reduced cognitive workload. Study 4 showed that the third dimension was not useful for search tasks in hierarchical data visualization in desktop VR. This study also showed that low spatial ability can significantly hinder performance in a virtual world, which implies that different visualization tools (or tools in other sensory modalities) will be needed for users with different cognitive or perceptual abilities. Finally, this thesis developed a methodology for studying visualization, which includes navigational logging and analysis, map projections of a 3D environment, search tasks with performance measures, and questionnaires to assess structure-learning ability. This methodology may provide a useful framework for subsequent research.

Reference List

1. Alexander, C., "A City is Not a Tree." In *Architectural Forum*, 1965. 122(1), 58-62, and 122(2), 58-61.
2. Alexander, C., *The Timeless Way of Building*. New York: Oxford University Press, 1979.
3. Anastasi, A., *Psychological Testing* (6th ed.). New York: Macmillan, 1988.
4. Andrews, K., "Visualizing Cyberspace: Information Visualization in the Harmony Internet Browser." In *Proceedings of IEEE Information Visualization '95*. IEEE, 1995, 97-104.
5. Andrews, K., Kappe, F., and Maurer, H., "Serving Information to the Web with Hyper-G." In *Proceedings of the Third International WWW Conference*. 1995, 919-926.
6. Bedersen, J., and Hollan, J., "Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics." In *Proceedings of the UIST '94 Conference*. ACM, 1994, 17-26.
7. Benedikt, M., "Cyberspace: Some Proposals." In Michael Benedikt, ed., *Cyberspace: First Steps*. Cambridge, Massachusetts: MIT Press, 1991, 119-224.
8. Berners-Lee, T., Cailliau, R., Luotonen, A., Nielsen, H.F., and Secret, A., "The World-Wide Web." In *Communications of the ACM*, 1994. 37(8), 76-82
9. Bush, V., "As We May Think." In *The Atlantic Monthly*, 1945. 176(1), 101-108.
10. Campagnoni, F.R., and Ehrlich, K., "Information Retrieval Using a Hypertext-Based Help System." In *ACM Transactions on Information Systems*, 1989. 7(3), 271-291.
11. Canter, D., "Cognitive Cartography." In *The Psychology of Place*. London: Architectural Press, 1977.
12. Canter, D., Rivers, R., and Storrs, G., "Characterizing User Navigation through Complex Data Structures." In *Behaviour and Information Technology*, 1995. 4(2), 93-102.
13. Card, S.K., Mackinlay, J.D., and Shneiderman, B., eds., *Readings in information visualization: Using vision to think*. San Francisco: Morgan Kaufmann, 1999.
14. Card, S.K., Robertson G.G., and Mackinlay, J.D., "The Information Visualizer, an Information Workspace." In *Proceedings of the CHI '91 Conference*. ACM, 1991, 181-188.
15. Catledge, L.D., and Pitkow, J.E., "Characterizing Browsing Strategies in the World Wide Web." In *Computer Networks and ISDN Systems 27*. Elsevier Science, 1995, 1065-1073.
16. Chalmers, M., "Using a landscape metaphor to represent a corpus of documents." In A.U. Frank and I. Campari, eds., *Spatial Information Theory: A Theoretical Basis for GIS: European Conference, COSIT '93*. Springer, 1993, 377-390.
17. Chalmers, M., Ingram, R., and Pfranger, C., "Adding Imageability Features to Information Displays." In *Proceedings of the UIST '96 Conference*. ACM, 1996, 33-39.
18. Chase, W.G., "Spatial Representations of Taxi Drivers." In R. Rogers and J.A. Sloboda, eds., *Acquisition of Symbolic Skills*. New York: Plenum, 1983, 391-405.
19. Chase, W.G., "Visual Information Processing." In K.R. Boff and J.E. Lincoln, eds., *Handbook of Human Perception and Performance*. New York: Wiley and Sons, 1986.
20. Chen, C., *Information Visualisation and Virtual Environments*. London: Springer, 1999.

21. Chen, C., and Czerwinski, M., "Spatial ability and visual navigation: An empirical study." In *The New Review of Hypermedia and Multimedia*, 1997. 3, 67-69.
22. Chen, C., Thomas, L., Cole, J., and Chennawasin, C., "Representing the semantics of virtual spaces." *IEEE Multimedia*, 1999. 6(2), 54-63.
23. Chi, E.H., Pirolli, P., and Pitkow, J., "The scent of a site: a system for analyzing and predicting information scent, usage, and usability of a Web site." In *Proceedings of the CHI '98 Conference*. ACM, 2000, 161-168.
24. Cronin, P., "Virtual field course evaluation." Position paper, Edinburgh Virtual Environment Centre, University of Edinburgh. 7 June 1998. http://www.iccs.informatics.ed.ac.uk/~paulus/Work/Vranded/posstatement.html
25. Czerwinski, M., van Dantzich, M., Robertson, G.G., and Hoffman, H., "The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D." In *Proceedings of HCI - INTERACT '99.* IFIP, 1999, 163-170.
26. Darken, R.P., and Sibert, J.L., "A Toolset for Navigation in Virtual Environments." In *Proceedings of the UIST '93 Conference*. ACM, 1993, 157-165.
27. Darken, R.P., and Sibert, J.L., "Wayfinding Strategies and Behaviors in Large Virtual Worlds." In *Proceedings of the CHI '96 Conference*. ACM, 1996, 142-149.
28. Dieberger, A., "Providing Spatial Navigation for the World Wide Web." In *Spatial Information Theory: A Theoretical Basis for GIS: International Conference, COSIT '95*. Berlin: Springer, 1995, 93-106.
29. Dieberger, A., "Browsing the WWW by interacting with a textual virtual environment -- A framework for experimenting with navigational metaphors." In *Proceedings of ACM Hypertext '96.* ACM, 1996, 170-179.
30. Dillon, A., McKnight, C., and Richardson, J., "Space--the Final Chapter or Why Physical Representations Are Not Semantic Intentions." In C. McKnight, A. Dillon, and J. Richardson, eds., *Hypertext: A Psychological Perspective*. New York: Ellis Horwood, 1993, 169-192.
31. Downs, R.M., and Stea, D., "Cognitive Maps and Spatial Behavior: Process and Products." In R.M. Downs and D. Stea, eds., *Image and Environment*. Chicago: Aldine, 1973, 8-26.
32. Elvins, T.T., Nadeau, D.R., and Kirsh, D., "Worldlets: 3D Thumbnails for Wayfinding in Virtual Environments." In *Proceedings of the UIST '97 Conference*. ACM, 1997, 21-30.
33. Elvins, T.T., Nadeau, D.R., Schul, R., and Kirsh, D., "Worldlets: 3D Thumbnails for 3D Browsing." In *Proceedings of the CHI '98 Conference*. ACM Press, 1998, 163-168.
34. Fairchild, K.M., Poltrock, S.E., and Furnas, G.W., "SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases." In Guindon, R., ed., *Cognitive Science and its Applications for Human-Computer Interactions*. Hillsdale, NJ: Erlbaum, 1988, 201-233.
35. Furnas, G.W., "Generalized Fish-eye Views." In *Proceedings of the CHI '86 Conference*. ACM, 1986, 16-23.

36. Furnas, G.W., “Effective View Navigation.” In *Proceedings of the CHI ’97 Conference.* ACM, 1997, 367-374.

37. Furnas, G.W., and Zacks, J., “Multitrees: Enriching and Reusing Hierarchical Structure.” In *Proceedings of the CHI '94 Conference*. ACM, 1994, 330-336.

38. Furnas, G.W., and Bederson, B.B., “Space-Scale Diagrams: Understanding Multiscale Interfaces.” In *Proceedings of the CHI ’95 Conference.* ACM, 1995, 234-241.

39. Golovchinsky, G., “Queries? Links? Is there a difference?” In *Proceedings of the CHI ’97 Conference*. ACM, 1997.

40. Golovchinsky, G., personal communication, 1997.

41. Hearst, M.A., and Karadi, C., “Cat-a-Cone: An Interactive Interface for Specifying Searches and Viewing Retrieval Results using a Large Category Hierarchy.” In *Proceedings of the ’97 SIGIR Conference*. ACM, 1997, 246-255.

42. Hemmje, M., Kunkel, C., and Willett, A., “LyberWorld -- A Visualization User Interface Supporting Fulltext Retrieval.” In *Proceedings of the ’94 SIGIR Conference*. Springer, 1994, 249-259.

43. Hendley, R.J., Drew, N.S., Wood, A.M., and Beale, R., “Narcissus: Visualising Information.” In *Proceedings of IEEE Visualization ‘95*. IEEE, 1995, 90-96.

44. Hutchins, E.L., Hollan, J.D., and Norman, D.A., “Direct Manipulation Interfaces.” In D.A. Norman and S.W. Draper, eds., *User Centered System Design: New Perspectives on Human-Computer Interaction*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1986, 87-124.

45. Ingram, R., and Benford, S., “Legibility Enhancement for Information Visualization.” In *Proceedings of IEEE Visualization ‘95*. 1995, 209-216.

46. Jul, S., and Furnas, G.W., “Navigation in electronic worlds: a CHI ’97 workshop.” In *SIGCHI Bulletin*, 1997. 29(4), 44-49.

47. Kaplan, N., and Moulthrop, S., “Where No Mind Has Gone Before: Ontological Design for Virtual Spaces.” In *European Conference on Hypermedia Technology 1994 Proceedings*. ACM, 1994, 206-216.

48. Keshkin, C., and Vogelmann, V., “Effective Visualization of Hierarchical Graphs with the Cityscape Metaphor.” In *Proceedings of the Workshop on “New paradigms in information visualization and manipulation.”* ACM, 1997, 52-57.

49. Kim, H., and Hirtle, S.C., “Spatial Metaphors and Disorientation in Hypertext Browsing.” In *Behaviour and Information Technology*, 1995. 14, 239-250.

50. Lakoff, G., and Johnson, M., *Metaphors We Live By*. Chicago: University of Chicago Press, 1980.

51. Lamping, J., Rao, R., and Pirolli, P., “A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies.” In *Proceedings of the CHI '95 Conference*. ACM, 1995, 401-408.

52. Likert, R., and Quasha, W.H., *Revised Minnesota Paper Form Board Test*. 1995: Harcourt, San Antonio, Texas.

53. Lynch, K., *Image of the City*, The. Cambridge, Massachusetts: The MIT Press, 1960.

54. Mackinlay, J.D., Rao, R., and Card, S.K., "An Organic User Interface for Searching Citation Links." In *Proceedings of the CHI '95 Conference*. ACM, 1995, 67-73.
55. McKnight, C., Dillon, A., and Richardson, J., *Hypertext in Context*. Cambridge, England: Cambridge University Press, 1991.
56. Mendelzon, A., "Visualizing the World Wide Web." White Paper. University of Toronto, 1996.
57. Modjeska, D., "Parallel Worlds: Spatial and Textual Representations of Information Structure." In *Workshop on Personalised and Social Navigation in Information Space*. Technical Report. Stockholm: Swedish Institute of Computer Science, 1998, 130-137.
58. Modjeska, D., "Designing for Navigation in Virtual Reality." In *Proceedings of HCI - INTERACT '99* (Vol. II). IFIP, 1999, 95-96.
59. Modjeska, D., and Waterworth, J., "Effects of 3D World Design on User Navigation and Search Performance." In *Proceedings of IV2000, the International Conference on Information Visualization*. IEEE, 2000.
60. Mukherjea, S., and Foley, J.D., "Showing the Context of Nodes in the World-Wide Web." In *Companion to the CHI '95 Conference*. ACM, 1995, 326-327.
61. Mukherjea, S., and Foley, J.D., "Visualizing the World-Wide Web with the Navigational View Builder." In *Computer Networks and ISDN Systems 27*. Elsevier Science, 1995, 1075-1087.
62. Mukherjea, S., Foley, J.D., and Hudson, S., "Visualizing Complex Hypermedia Networks through Multiple Hierarchical Views." In *Proceedings of the CHI '95 Conference*. ACM, 1995, 331-337.
63. Neill, J., "The AOL Virtual Reality Resource Center FAQ (Frequently Asked Questions)." American Online, 1 November 1994. http://www.hitl.washington.edu/projects/knowledge_base/virtual-worlds/scivw-ftp/other/FAQ-AOL-VR
64. Nielsen, J., *Multimedia and Hypertext: The Internet and Beyond*. Boston: AP Professional, 1995.
65. Passini, R., *Wayfinding in Architecture*. New York: Van Nostrand Reinhold, 1984.
66. Perlin, K., and Fox, D., "Pad: An Alternative Approach to the Computer Interface." In *Proceedings of the SIGGRAPH '93 Conference*. ACM, 1993, 57-64.
67. Pirolli, P., "Computational Models of Information Scent-Following in a Very Large Browsable Text Collection." In *Proceedings of the CHI '97 Conference.* ACM: New York, 1997, 3-10.
68. Pirolli, P., and Card, S., "Information Foraging in Information Access Environments." In *Proceedings of the CHI '95 Conference*. ACM, 1995, 51-58.
69. Pirolli, P., and Card, S., "Information Foraging." In *Psychological Review*, 1999. 106(4), 643-675.
70. Rennison, E., "Galaxy of News: An Approach to Visualizing and Understanding Expansive News Landscapes." In *Proceedings of the UIST '94 Conference*. ACM, 1994, 3-12.
71. Rennison, E., and Strausfeld, L., "The Millennium Project." In A.U. Frank and W. Kuhn, eds., *Spatial Information Theory: A Theoretical Basis for GIS: International Conference, COSIT '95*. Berlin: Springer, 1995, 69-91.
72. Robertson, G.G., Mackinlay, J.D., and Card, S.K., "Cone Trees: Animated 3D Visualizations of Hierarchical Information." In *Proceedings of the CHI '91 Conference*. ACM, 1991, 189-194.

73. Rojdestvenski, I., Modjeska, D., Pettersson, F., Rojdestvenskaia, M., and Gustafsson, P., "Sequence World: A Genetics Database in Virtual Reality." In *Proceedings of IV2000, the International Conference on Information Visualisation*. IEEE, 2000.
74. Sarkar, M., and Brown, M.H., "Graphical Fisheye Views of Graphs." In *Proceedings of the CHI '92 Conference*. ACM, 1992, 83-91.
75. Satalich, G.A., "Navigation and Wayfinding in Virtual Reality: Finding the Proper Tools and Cues to Enhance Navigational Awareness." M.S. Thesis. University of Washington, 1997.
76. Shum, S.B., "Real and virtual spaces: mapping from spatial cognition to hypertext." In *Hypermedia*, 1990. 2(2), 133-158.
77. Stanney, K.M., Mourant, R.R., Kennedy, R.S., "Human Factors Issues in Virtual Environments: A Review of the Literature." In *Presence*, 1998. 7(4), 327-351.
78. Stevens, A., and Coupe, P., "Distortions in judged spatial relations." In *Cognitive Psychology*, 1978. 10, 422-437.
79. Tauscher, L., and Greenberg, S., "Revisitation Patterns in World Wide Web Navigation." In *Proceedings of the CHI '97 Conference.* ACM, 1997, 399-406.
80. Thorndyke, P.W., and Goldin, S.E., "Spatial Orientation: Theory, Research, and Application." In *Spatial Learning and Reasoning Skill*. Plenum Press: New York, 1983, 195-217.
81. Thorndyke, P.W., and Hayes-Roth, B., "Differences in Spatial Knowledge Acquired from Maps and Navigation." In *Cognitive Psychology*, 1982. 14, 560-589.
82. Tompa, F.W., "A data model for flexible hypertext database systems." In *ACM Transactions on Information Systems*, 1989. 7(1), 85-100.
83. Tweedie, L., "Characterizing Interactive Externalizations." In *Proceedings of the CHI '97 Conference.* ACM, 1997, 375-382.
84. Valdez, F., Chignell, M., and Glenn, B., "Browsing Models for Hypermedia Databases." In *Proceedings of the Human Factors Society, 32nd Annual Meeting*. 1988, 318-322.
85. Vinson, N.G., "Design Guidelines for Landmarks to Support Navigation in Virtual Environments." In *Proceedings of the CHI '99 Conference*. ACM, 1999, 278-285.
86. Waller, D.A. *An assessment of individual differences in spatial knowledge of real and virtual environments*. Ph.D. thesis. University of Washington, 1999.
87. Waterworth, J.A. *The Influence of Variations in Cognitive Processing on the Perception of Time*. Ph.D. thesis. University of Hertfordshire, UK, 1983.
88. Waterworth, J A, "A pattern of islands: exploring public information space in a private vehicle." In P. Brusilovsky, P. Kommers and N. Streitz, eds., *Multimedia, Hypermedia and Virtual Reality: Models, Systems, and Applications*. Berlin: Springer, 1996, 266-279.

89. Waterworth, J.A., “Personal Spaces: 3D Spatial Worlds for Information Exploration, Organisation and Communication.” In R. Earnshaw and J. Vince, eds., *'The Internet in 3D: Information, Images, and Interaction*. New York: Academic Press, 1997.
90. Waterworth, J.A., and Chignell, M.H., “A Model for Information Exploration.” In *Hypermedia*. 1991. 3(1), 35-58.
91. Wickens, C.D., “Spatial Perception and Cognition and the Display of Spatial Information.” In *Engineering Psychology and Human Performance*. New York: Harper Collins, 1992, 116-166.
92. Wurman, R.S., *Information Anxiety*. New York: Bantam, Doubleday, Dell, 1989.

Appendix A

Study 1 – Design, Script, and Questionnaire

Subject	Block	Des	Targets
1	1	Night	2
1	2	Dusk	3
1	3	Day	1
2	1	Day	3
2	2	Night	1
2	3	Dusk	2
3	1	Dusk	1
3	2	Day	2
3	3	Night	3
4	1	Dusk	2
4	2	Night	3
4	3	Day	1
5	1	Night	1
5	2	Day	2
5	3	Dusk	3
6	1	Day	3
6	2	Dusk	1
6	3	Night	2
7	1	Night	2
7	2	Dusk	3
7	3	Day	1
8	1	Day	3
8	2	Night	1
8	3	Dusk	2
9	1	Dusk	1
9	2	Day	2
9	3	Night	3
10	1	Dusk	2
10	2	Night	3
10	3	Day	1
11	1	Night	1
11	2	Day	2
11	3	Dusk	3
12	1	Day	3
12	2	Dusk	1
12	3	Night	2

Table A.1. Experimental design for Study 1.

Script for Study 1
Parallel Worlds Project
Umeå University, 1999

0) **Preparation**
   a) 2 blank notepad pages
   b) 6 blank sketchpad pages
   c) 3+ colored pens
   d) C60+ cassette tape rewound in recorder
1) **Introduction**
   a) Experimenter welcomes subject
      i) Experimenter introduces self – name, title, affiliation(s)
   b) Experimenter gives brief explanation of study
      i) About observing people using virtual reality
      ii) Hardware - a Silicon Graphics' Onyx2 Infinite Reality computer.
      iii) Software - Netscape Communicator WWW browser
         (1) With CosmoPlayer plug-in to view VRML files
      iv) Study data will be reported anonymously in my Ph.D. thesis and perhaps an academic article
         (1) Copy available on request
2) **Computer Training and Practice**
   a) Experimenter explains Cosmo Player UI
      i) Easiest method - Shift+click on something to fly towards it
      ii) With left button down, mouse controls forward/back and right/left spin
      iii) With right button down, mouse controls up/down and right/left slide
      iv) Shift key accelerates flying
      v) With Ctrl key + left button down, mouse controls looking in all directions
      vi) "Horizontal" button returns user to world origin
   b) Subject is allowed 5-15 minutes practice in virtual Aztec temple compound
      i) Local file ~/public_html/Aztec/index.html
      ii) Subject decides when ready to proceed with session
3) **Task explanation**
   a) Subject will explore three virtual worlds
      i) Provide feedback after each world
         (1) These verbal comments will be audio recorded – Is this OK?
      ii) Provide feedback at end of study session
   b) In each world, there are a number of targets to be found
      i) When a target is found:
         (1) Click on target with the mouse
      ii) The mouse click will play a unique sound
         (1) Demonstrate clicking on target to play sound
         (2) Local file ~/public_html/PW/sample.wrl
      iii) Subject's goal is to find and write down as many sounds as possible in time available
   c) A set of targets will be given for each world, to help with finding sounds
      i) Targets will be given on paper cards, in sequence
      ii) Each card will show a hint (at top) and a target name (at bottom).
      iii) If a target is hard to find, here are some suggestions:
         (1) as appropriate after four minutes seeking one target
         (2) More details are usually visible close to an object.
         (3) If necessary, spin around to survey the environment
         (4) If necessary, seek elsewhere in world
         (5) If necessary, skip to next target

   d) When a target is found, write down card number and sound description
      i) Keep a simple list on a notepad sheet
   e) Questions are OK anytime; breaks OK between virtual worlds

4) **Per-World Tasks (Repeat three times)**
   a) Subject clicks on a virtual world name
      i) Experimenter has counterbalanced the world order per subject
      ii) Local file ~/public_html/PW/index.html
      iii) Subject waits about two minutes for world to appear on screen
      iv) Subject explores world for about a minute
   b) Experimenter picks up a set of target cards
      i) Experimenter has counterbalanced the set order per subject, counterbalanced the set assignment to worlds, randomized the target assignments to sets, and randomized the target orders within each set
   c) Subject locates as many sounds as possible
      i) Experimenter notes hunt's start time, and when each target is found.
      ii) Experimenter quietly limits hunting time to 20 minutes
   d) Subject clicks on browser's back button
   e) Subject speaks world name into tape recorder
   f) Subject briefly narrates experience in world into tape recorder (5 min.)
      i) English or Swedish is OK
   g) Subject sketches world in as much detail as possible on A3 pad (5 min.)
      i) Subject may use as many pages as desired

5) **Summary Questionnaire**

6) **Conclusion**
   a) Experimenter thanks subject for participation in study.

Summary Questionnaire
Parallel Worlds Study
Umeå University
March 1999

1) Please rate the worlds according to size.

   a) Color world ______
   b) Gray world ______
   c) Text world ______

Very small	= 1
Small	= 2
Medium	= 3
Large	= 4
Very large	= 5

2) Please rate the worlds according to duration of your sessions.

   a) Color world ______
   b) Gray world ______
   c) Text world ______

Very short	= 1
Short	= 2
Medium	= 3
Long	= 4
Very long	= 5

3) Please rate the worlds according to sense of presence – “You are there.”

   a) Color world ______
   b) Gray world ______
   c) Text world ______

Very weak	= 1
Weak	= 2
Medium	= 3
Strong	= 4
Very strong	= 5

4) Please rate the worlds according to ease of use.

   a) Color world ______
   b) Gray world ______
   c) Text world ______

Very easy	= 1
Easy	= 2
Medium	= 3
Hard	= 4
Very hard	= 5

5) Please rate the worlds according to personal preference.

   a) Color world ______
   b) Gray world ______
   c) Text world ______

Strongly disliked	= 1
Disliked	= 2
Neutral	= 3
Liked	= 4
Strongly liked	= 5

6) What is your age? ______
7) What are your occupation and field?
8) What languages do you know fluently, e.g., English, Swedish?

9) How often do you use computers? (Please check one option.)
   a) Four hours per day or more
   b) An hour per day
   c) An hour per week
   d) An hour per month
   e) An hour per year or less

10) How often do you play computer or video games? (Please check one option.)
    a) An hour per day or more
    b) An hour per week
    c) An hour per month
    d) An hour per year or less

11) How often do you use the Yahoo Web index? (Please check one option.)
    a) Once a day or more
    b) Once a week
    c) Once a month
    d) Once a year or less

12) Do you have any general comments on the worlds or your experiences during the study?

Appendix B

Study 1 – Data Summary

		SUBJ	BLOCK	DESIGN	TARGETS	SIZE	DURATION	PRESENCE	EASE	ENJOY
1		1	1	night	1	2	5	3	3.00	3
2		1	2	dusk	3	5	4	2	2.00	2
3		1	3	day	6	3	2	4	4.00	5
4		2	1	day	8	3	4	3	3.00	3
5		2	2	night	8	3	2	3	4.00	4
6		2	3	dusk	1	3	3	1	1.00	1
7		3	1	dusk	6	1	1	1	2.00	3
8		3	2	day	7	3	3	3	5.00	5
9		3	3	night	5	5	5	4	1.00	1
10		4	1	night	3	5	4	2	3.00	3
11		4	2	day	2	2	2	5	5.00	4
12		4	3	dusk	5	2	3	5	4.00	2
13		5	1	dusk	4	2	4	3	3.00	2
14		5	2	night	4	5	3	1	2.00	3
15		5	3	day	6	4	2	5	5.00	5
16		6	1	day	1	2	2	2	2.00	2
17		6	2	dusk	2	3	4	3	3.00	3
18		6	3	night	4	4	3	4	4.00	4
19		7	1	night	3	3	3	3	3.00	3
20		7	2	dusk	0	2	3	2	1.00	1
21		7	3	day	9	4	4	4	4.00	4
22		8	1	day	3	4	4	3	2.00	2
23		8	2	night	3	3	3	2	3.00	4
24		8	3	dusk	1	5	5	3	1.00	2
25		9	1	dusk	7	3	3	3	3.00	3
26		9	2	day	7	3	3	4	3.00	4
27		9	3	night	6	4	3	1	2.00	2
28		10	1	night	5	5	3	2	3.00	2
29		10	2	day	8	2	3	4	4.00	4
30		10	3	dusk	2	3	2	3	3.00	3
31		11	1	dusk	3	3	1	3	2.00	2
32		11	2	night	4	5	3	4	3.00	1
33		11	3	day	4	2	2	5	2.00	5
34		12	1	day	1	3	5	4	2.00	4
35		12	2	dusk	4	3	4	3	1.00	2
36		12	3	night	8	2	2	1	3.00	3
Total	N	36	36	36	36	36	36	36	36	36

Table B.1. Case summaries for Study 1.

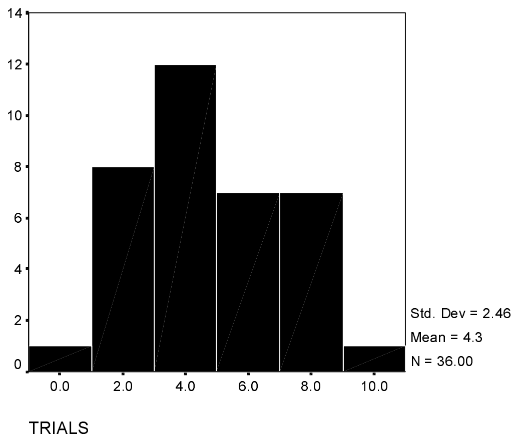

Figure B.1. Histogram of Targets measure in Study 1.

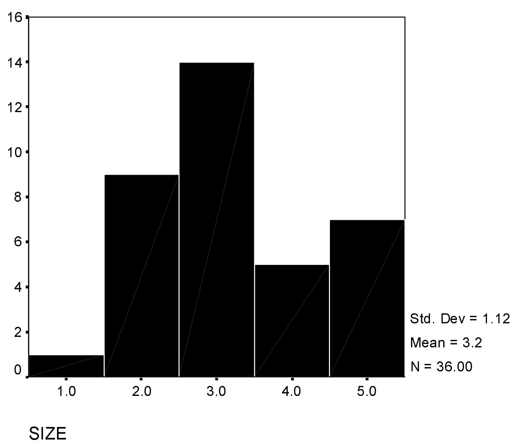

Figure B.2. Histogram of Size measure in Study 1.

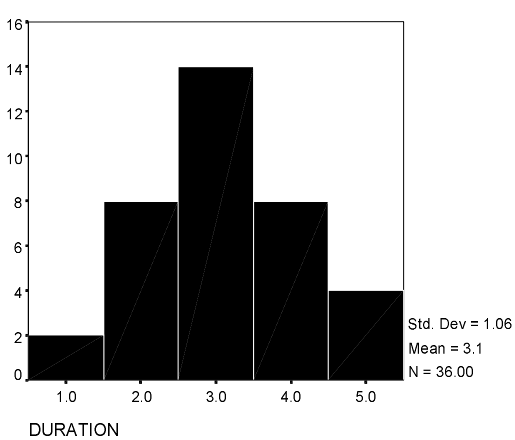

Figure B.3. Histogram of Duration measure in Study 1.

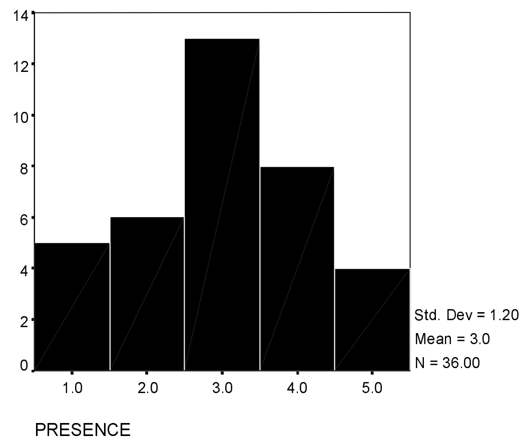

Figure B.4. Histogram of Presence measure in Study 1.

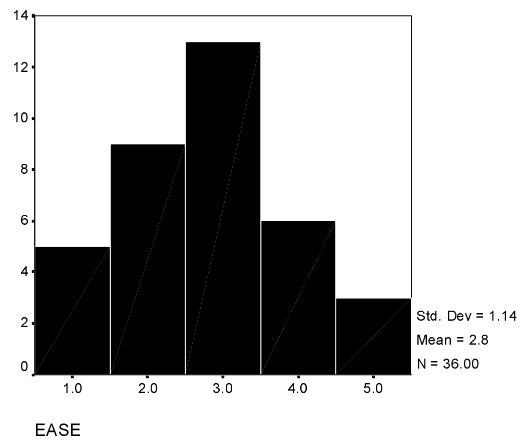

Figure B.5. Histogram of Ease measure in Study 1.

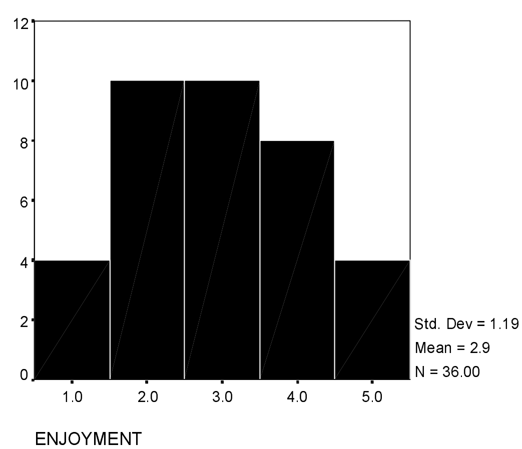

Figure B.6. Histogram of Enjoyment measure in Study 1.

Appendix C

Study 1 - Sample Drawings

Figure C.1. Study 1 – Sample Drawing 1.

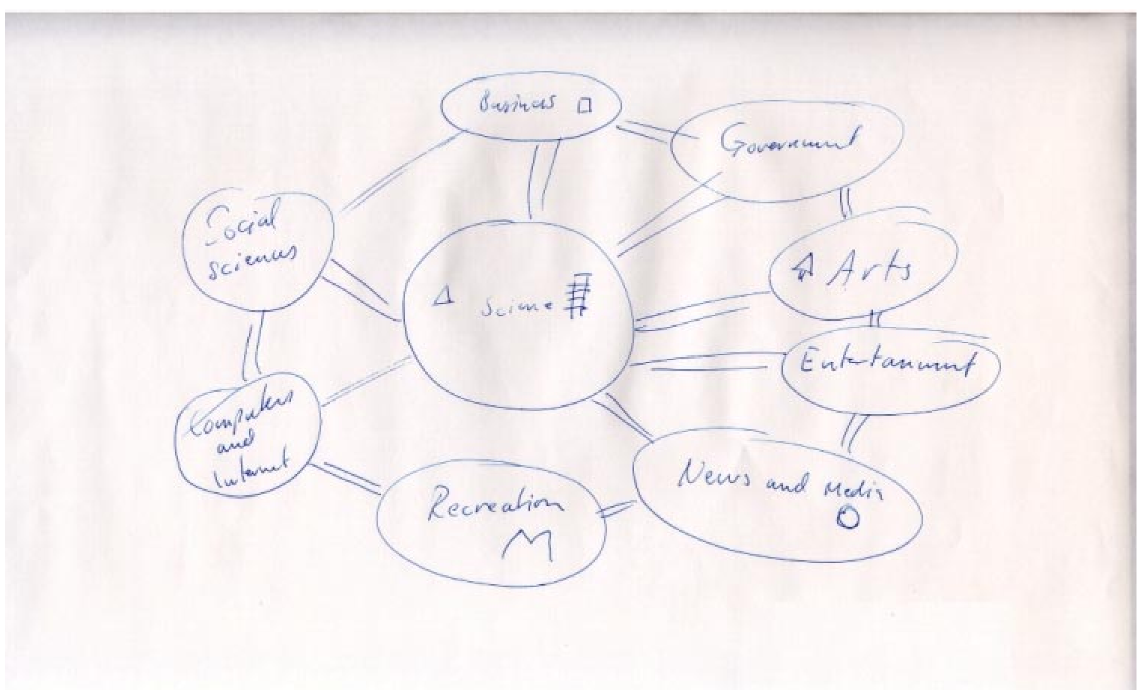

Figure C.2. Study 1 – Sample Drawing 2.

Appendix D

Study 2 – Design, Script, and Questionnaires

Subject	Block	Design	Data
1	1	vr	Work
1	2	ht	Work
1	3	vr	Play
1	4	ht	Play
2	1	ht	Play
2	2	vr	Play
2	3	ht	Work
2	4	vr	Work
3	1	vr	Play
3	2	ht	Play
3	3	vr	Work
3	4	ht	Work
4	1	ht	Work
4	2	vr	Work
4	3	ht	Play
4	4	vr	Play
5	1	ht	Work
5	2	ht	Play
5	3	vr	Work
5	4	vr	Play
6	1	vr	Work
6	2	vr	Play
6	3	ht	Work
6	4	ht	Play
7	1	ht	Play
7	2	ht	Work
7	3	vr	Play
7	4	vr	Work
8	1	vr	Play
8	2	vr	Work
8	3	ht	Play
8	4	ht	Work

Subject	Block	Design	Data
9	1	vr	Work
9	2	ht	Work
9	3	vr	Play
9	4	ht	Play
10	1	ht	Play
10	2	vr	Play
10	3	ht	Work
10	4	vr	Work
11	1	vr	Play
11	2	ht	Play
11	3	vr	Work
11	4	ht	Work
12	1	ht	Work
12	2	vr	Work
12	3	ht	Play
12	4	vr	Play
13	1	ht	Work
13	2	ht	Play
13	3	vr	Work
13	4	vr	Play
14	1	vr	Work
14	2	vr	Play
14	3	ht	Work
14	4	ht	Play
15	1	ht	Play
15	2	ht	Work
15	3	vr	Play
15	4	vr	Work
16	1	vr	Play
16	2	vr	Work
16	3	ht	Play
16	4	ht	Work

Table D.1. Experimental design for Study 2.

Parallel Worlds Project
Script for Study 2
Umeå University, May 1999

1) **Preparation**
   a) Restart computer
   b) 2 blank notepad pages
   c) Questionnaire with five parts, labeled and ordered per subject
   d) Close blinds and pull cords to turn down lights
   e) Turn off email and ICQ
2) **Introduction**
   a) Welcome
      i) Experimenter thanks subject for participation
      ii) Experimenter and subject introduce selves, if not acquainted already
   b) Explanation of experiment
      i) For Ph.D. research about people's experience of electronic worlds
      ii) For experiments, VR and hypertext worlds developed at Umeå University
      iii) We will see VR in CosmoPlayer plug-in for Netscape Communicator
      iv) We will see hypertext in Netscape Communicator
      v) H/w - Dell 400 MHz Pentium PC w/ Diamond Fire GL 4000 grafix card
3) **Experiment instructions**
   a) Total time about 90 minutes
      i) 10 minutes VR training time
      ii) Exploring 4 electronic worlds + short questionnaire for each one
      iii) Short questionnaire at the end, + general –interest questions afterwards
   b) In each electronic world
      i) Minute of free exploration to begin
      ii) Hunt per list of targets to find as many as possible (in order)
         (1) Each list item shows a hint, plus the target's exact name in italics
         (2) There are more targets listed than anybody is expected to find
      iii) When a target is found
         (1) In VR, mouse cursor will change to ring over target
         (2) Mouse click on target to play a sound, then close sound window
         (3) *Then* press "T" key to log target in computer
   c) Hints for finding targets (repeat every four minutes as necessary)
         (1) All buildings contain information items
         (2) Approach an object and/or look around to see more information
         (3) Survey the environment frequently by looking around
         (4) All targets present – seek elsewhere if current location seems wrong
         (5) Subject may skip an item & move on. Tell experimenter & press 'S' key
   d) Questions are OK anytime, and breaks are allowed between worlds

4) **Computer Training and Practice**
   a) Experimenter explains and demos VR UI
      i) Select SGI's Aztec demo world from bookmarks (1-min. download)
      ii) Flying horizontally –With left button down, drag mouse forward, back, right, left in VR window
         (1) Stretching "rubber band" line longer increases speed, but loses control
      iii) Zoom flying – Shift key plus (a) mouse drag or (b) click on object goal
      iv) Flying vertically – With right button down, drag mouse forward, back
      v) Returning home – click mouse on "Home" item at bottom left of screen
      vi) After long mouse drags, the computer may delay 1-5 seconds to catch up
   b) Subject practices VR UI up to 10 minutes in SIG sample world
      i) Select "Aztec" item on bookmarks menu
      ii) Wait about one minute for download
      iii) Drag horizontal frame boundary up to window top
   c) Experimenter verifies that subject has WWW experience or explains
      i) Mention menu-path bar at top of page in Flat_World_all.html
5) **Experimenter's tasks per world**
   a) Experimenter sets up Netscape
      i) Open correct electronic world from bookmarks list (up to 1 min.)
      ii) Hide taskbars
   b) Experimenter sets up Spy
      i) Verify that previous log file has been renamed appropriately
      ii) Choose Netscape client area *and press OK*
   c) Subject explorers electronic world
      i) Subject explores freely one minute, then hits "Home", and "H" key for log
      ii) Experimenter gives subject a target list
      iii) Subject hunts for targets (Warn after 14 minutes, stop after 15.)
      iv) When a target is found or discarded, experimenter notes on copy of list
   d) Follow-up for exposure
      i) Experimenter and subject swap chairs
      ii) Subject gets and fills in 1st questionnaire page about world
      iii) Subject *then* gets and fills in 2nd questionnaire page about world
      iv) Experimenter stops Spy logging, exits log window, and renames log file
         (1) E.g., "subj3_world2.txt"
   e) Repeat (4) three times
6) **Conclusion**
   a) Subject completes summary questionnaire
   b) Experimenter gives subject 100 crowns and requests signature on form
   c) Experimenter thanks subject for participating in experiment

"Work" Targets 1

1. Real estate companies in *Switzerland*
2. *Art* in Egyptian archaeology
3. *Comet Halley*
4. The American *National Museum of Natural History*
5. *World Wide Web Consortium & xNet Consortium*
6. Computer companies making *Embedded Control* hardware
7. **Financial services companies in the *Netherlands***
8. Home-and-garden companies that do *Interior Plantscaping*
9. Companies that provide *Arts Therapy*
10. *Trade Magazines* about security companies
11. Business opportunities for *Restaurants*
12. The linguist *de Saussure Ferdinand 1857 1913*
13. Construction companies that make *Storm and Security Shields*
14. The insurance company *Protective Life Corporation*
15. The *Genome Projects* in human genetics
16. *National Information Infrastructure US & IT Policy OnRamp*
17. Real estate companies in *Bahamas The*
18. The study of hurricanes in *1996*
19. Construction companies that make *Steel Framed Homes*
20. The credit-card company *NOVUS*
21. The British *Chartered Institute of Transport*
22. Financial services companies in *Lithuania*
23. Financial services companies in *Costa Rica*
24. *Directories* of companies that make guns
25. An American law called the *Employment Non Discrimination Act ENDA*

“Work” Targets 2

1. The study of *Seismic Safety* (in earthquakes)
2. *Treaties Pacts and Agreements* on intellectual property
3. The *Ridesharing* (carpooling) business
4. Economics Organizations for the *Student*
5. Internet Radio *Stations*
6. The study of *Mollusks* (e.g., snails, clams)
7. *Personal Accounts* by American veterans of the Vietnam War
8. Real estate companies in *Honduras*
9. *Indices* for seismology (the study of earthquakes)
10. The study of *Insects*
11. Financial services companies in *Taiwan*
12. Financial services companies in *Egypt*
13. The American *National Institutes of Health NIH*
14. Real estate companies in *Canada*
15. *Development Banks* (e.g., investors in industrializing nations)
16. The astronomer *Brahe Tycho 1546 1601*
17. Real estate companies in *Malaysia*
18. Construction companies making *Steel Frame Packages*
19. Computer companies that make *Mainframes*
20. Home building companies that make a *Timber Frame*
21. *Institutes* for transportation engineers
22. Construction companies that make *Wallboard*
23. *Classifieds* (Ads) for companies in card collecting (e.g., baseball cards)
24. Legal *Products & SoCoOL Commentary on O.J. Simpson*
25. Real estate companies in *Sri Lanka*

“Work” Questionnaire 1

Please classify the following items on as many levels as you can, according to the electronic world you have just seen. It’s OK to guess. Here is an example for *Sailing*:

Recreation –> Outdoors -> Sailing

1. Financial services companies in *Guyana*

2. *United States Information Agency USIA*

3. Public policy studies by *World Affairs Council & Progress and Freedom Foundation*

4. Studies of the *1989 Loma Prieta Earthquake*

5. Internet *Fraud*

(continued on next page)

"Work" Questionnaire 1 (cont.)

1) Please rate the information categories according to size.

a) Business ______
b) Computers & Internet ______
c) Government ______
d) Science ______
e) Social Science ______

Very small	= 1
Small	= 2
Medium	= 3
Large	= 4
Very large	= 5

2) Please estimate the length of your session: ______ minutes

“Work” Questionnaire 2

Please classify the following items on as many levels as you can, according to the electronic world you have just seen. It’s OK to guess. Here is an example for *Sailing*:

Recreation –> Outdoors -> Sailing

1. *Commercial Website Directories* of Heating, Ventilation, and Air Conditioning (HVAC) Companies

2. The ethics of *Human Subjects* for biomedical research

3. *Americans for Constitutional Action & Flat Tax* reform

4. Internet Radio *Shows*

5. Urban Studies by the *American Planning Association*

(continued on next page)

“Work” Questionnaire 2 (cont.)

1) Please rate the information categories according to size.

f) Business	______	Very small	= 1
g) Computers & Internet	______	Small	= 2
h) Government	______	Medium	= 3
i) Science	______	Large	= 4
j) Social Science	______	Very large	= 5

2) Please estimate the length of your session: ______ minutes.

"Play" Targets 1

1. TV cartoon *Underdog*
2. The history of *Ancient Art*
3. Pop group *Spice Girls*
4. *Museums and Memorials* of military history
5. The Viking *Ericsson Leif 970 1020*
6. Baroque music *Ensembles*
7. Rock musician *Young Neil*
8. A Marvel comic book called *Thunderbolts*
9. Pop musician *Denver John 1943 1997*
10. *Violinists*
11. Internet Relay Chat (*IRC*) about music
12. The children's TV show *Clarissa Explains It All*
13. Philosopher *Foucault Michel 1926 1984*
14. TV police show *Homicide*
15. Musicians in *Luxembourg*
16. *Women s History Month*
17. Humorous music *Music That Sucks*
18. TV sketch comedy *Roundhouse*
19. *Dixieland* jazz musicians
20. History of the *South Africa* region
21. *Dub* (Reggae) musicans
22. Rock/pop musician *Stefani Gwen*
23. Ancient Greek *Alexander the Great*
24. Musicians in the *Ska* genre
25. *Masters* of photography

"Play" Targets 2

1. The comic book *Tick The*
2. TV soap opera *As The World Turns*
3. History of the *Asia* region
4. The comic book *Spawn*
5. The TV cartoon *Pinky and the Brain*
6. TV action show *Viper*
7. Maritime (shipping) history of the *Titanic*
8. *Native American* musicians
9. Music *Recording Equipment*
10. Funny *Song Parodies*
11. *Maps* of the Medieval Crusades
12. The DC comic book *Justice Society of America*
13. TV dramas *Wonder Years The & Zorro*
14. *Reader Reviews* of music
15. The Marvel comic book *Spider Man*
16. The animated cartoon *Speed Racer*
17. The *Impressionism* movement in art
18. Philosopher *Heidegger Martin 1889 1976*
19. The literary genre of *Storytelling*
20. *Indices* of Christian musicians
21. TV police show *21 Jump Street*
22. Children's TV show *Blue s Clues*
23. The film *Gone With The Wind*
24. Recipes for *Soups*
25. *Organizations* in 19th-century American history

“Play” Questionnaire 1

Please classify the following items on as many levels as you can, according to the electronic world you have just seen. It’s OK to guess. Here is an example for *Sailing*:

Recreation –> Outdoors -> Sailing

1. Rock/pop musician *Martin Dean*

2. American *Historical Societies*

3. TV dramas *Cape The & Avonlea*

4. Folk musician *Buckley Jeff 1966 1997*

5. The *Bibliographical Society* on the history of books and printing

(continued on next page)

"Play" Questionnaire 1 (cont.)

1) Please rate the information categories according to size.

a) Arts ______
b) Entertainment ______
c) News & media ______

Very small	= 1
Small	= 2
Medium	= 3
Large	= 4
Very large	= 5

2) Please estimate the length of your session: ______ minutes

“Play” Questionnaire 2

Please classify the following items on as many levels as you can, according to the electronic world you have just seen. It’s OK to guess. Here is an example for *Sailing*:

Recreation –> Outdoors -> Sailing

1. TV *Cooking Shows*

2. Philosopher *Baudrillard Jean*

3. The comic book *Shi*

4. *Weekly* food recipes

5. *Disc Jockeys* in the reggae genre

(continued on next page)

“Play” Questionnaire 2 (cont.)

1) Please rate the information categories according to size.

a)	Arts	______
b)	Entertainment	______
c)	News & Media	______

Very small	= 1
Small	= 2
Medium	= 3
Large	= 4
Very large	= 5

2) Please estimate the length of your session: ______ minutes.

Summary Questionnaire
Parallel Worlds Study 2
Umeå University
May 1999

1) Please rate the worlds according to size.

a) Virtual reality world ______
b) Hypertext world ______

Very small	= 1
Small	= 2
Medium	= 3
Large	= 4
Very large	= 5

2) Please rate the worlds according to sense of presence – "You are there."

a) Virtual reality world ______
b) Hypertext world ______

Very weak	= 1
Weak	= 2
Medium	= 3
Strong	= 4
Very strong	= 5

3) Please rate the worlds according to ease of use.

a) Virtual reality world ______
b) Hypertext world ______

Very hard	= 1
Hard	= 2
Medium	= 3
Easy	= 4
Very easy	= 5

4) Please rate the worlds according to efficiency.

a) Virtual reality world ______
b) Hypertext world ______

Very inefficient	= 1
Inefficient	= 2
Neutral	= 3
Efficient	= 4
Very efficient	= 5

5) Please rate the worlds according to enjoyment.

a) Virtual reality world ______
b) Hypertext world ______

Very unpleasant	= 1
Unpleasant	= 2
Neutral	= 3
Pleasant	= 4
Very pleasant	= 5

Appendix E

Study 2 – Data Summary

		SUBJ	BLOCK	DES	DATA	TAR-GETS	DURATION	SIZE	PRESENCE	EASE	EFFIC	ENJOY	STRUCT TEST
1		1	1	vr	2	3	5	4	3	4	4	2	.60
2		1	2	ht	2	18	10	2	4	2	2	4	.50
3		1	3	vr	1	4	10	4	3	4	4	2	.60
4		1	4	ht	1	14	15	2	4	2	2	4	.90
5		2	1	ht	1	14	20	3	2	3	4	3	.40
6		2	2	vr	1	2	20	4	4	2	4	5	.70
7		2	3	ht	2	11	20	3	2	3	4	3	.60
8		2	4	vr	2	6	20	4	4	2	4	5	.40
9		3	1	vr	1	0	10	5	2	1	1	4	.50
10		3	2	ht	1	6	8	4	4	4	4	3	.60
11		3	3	vr	2	3	8	5	2	1	1	4	.50
12		3	4	ht	2	8	9	4	4	4	4	3	.70
13		4	1	ht	2	10	15	3	2	4	4	3	.30
14		4	2	vr	2	3	15	3	4	2	2	2	.40
15		4	3	ht	1	14	15	3	2	4	4	3	.70
16		4	4	vr	1	7	15	3	4	2	2	2	.90
17		5	1	ht	2	11	15	4	2	5	4	3	.70
18		5	2	ht	1	18	20	4	2	5	4	3	.90
19		5	3	vr	2	2	20	5	4	2	2	4	.60
20		5	4	vr	1	5	20	5	4	2	2	4	1.00
21		6	1	vr	2	3	10	4	3	3	3	4	.40
22		6	2	vr	1	1	10	4	3	3	3	4	.70
23		6	3	ht	2	11	8	5	4	4	4	3	.30
24		6	4	ht	1	16	12	5	4	4	4	3	.90
25		7	1	ht	2	11	20	4	3	4	4	3	.70
26		7	2	ht	1	13	20	4	3	4	4	3	.60
27		7	3	vr	2	2	15	4	4	2	2	5	.20
28		7	4	vr	1	4	30	4	4	2	2	5	.30
29		8	1	vr	2	0	4	5	3	1	1	2	.50
30		8	2	vr	1	2	10	5	3	1	1	2	.50
31		8	3	ht	2	21	10	4	4	4	5	4	1.00
32		8	4	ht	1	15	10	4	4	4	5	4	.80
33		9	1	vr	2	2	10	4	3	2	2	4	.30
34		9	2	ht	2	7	20	2	5	5	5	1	.40
35		9	3	vr	1	7	25	4	3	2	2	4	.20
36		9	4	ht	1	8	28	2	5	5	5	1	.00
37		10	1	ht	1	6	10	3	4	5	5	2	.60
38		10	2	vr	1	1	8	5	2	3	2	4	.40
39		10	3	ht	2	11	16	3	4	5	5	2	.20
40		10	4	vr	2	4	15	5	2	3	2	4	.40
41		12	1	ht	2	7	15	4	4	4	3	3	.50
42		12	2	vr	2	2	10	3	2	2	2	4	.70
43		12	3	ht	1	10	15	4	4	4	3	3	.70
44		12	4	vr	1	5	10	3	2	2	2	4	.70
45		13	1	ht	2	7	15	3	2	4	4	2	.30
46		13	2	ht	1	4	10	3	2	4	4	2	.30
47		13	3	vr	2	0	10	3	4	1	2	4	.10
48		13	4	vr	1	2	10	3	4	1	2	4	.40
49		14	1	vr	2	4	15	4	4	2	2	4	.20
50		14	2	vr	1	3	15	4	4	2	2	4	.20
51		14	3	ht	2	11	15	4	2	4	4	3	.60
52		14	4	ht	1	18	15	4	2	4	4	3	.90
53		15	1	ht	2	6	15	4	2	3	3	3	.20
54		15	2	ht	1	11	15	4	2	3	3	3	.10
55		15	3	vr	2	2	15	5	3	2	2	5	.00
56		15	4	vr	1	2	15	5	3	2	2	5	.20
57		16	1	vr	2	1	10	5	5	1	1	5	.00
58		16	2	vr	1	4	10	5	5	1	1	5	.20
59		16	3	ht	2	10	15	2	1	5	5	1	.60
60		16	4	ht	1	13	20	2	1	5	5	1	.40
Total	N	60	60	60	60	60	60	60	60	60	60	60	60

Table E.1. Case summaries for Study 2.

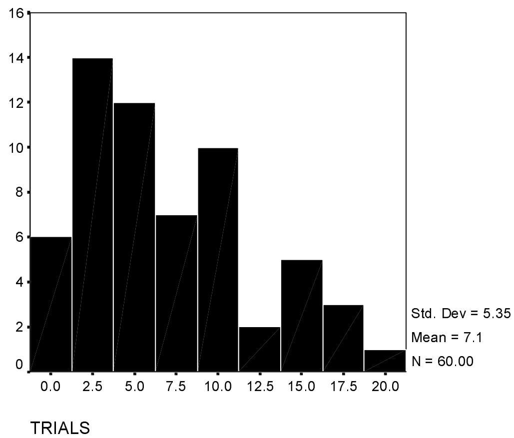

Figure E.1. Histogram of Targets measure in Study 2.

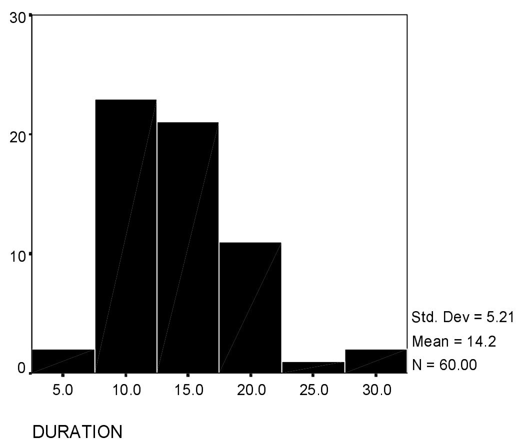

Figure E.2. Histogram of Duration measure in Study 2.

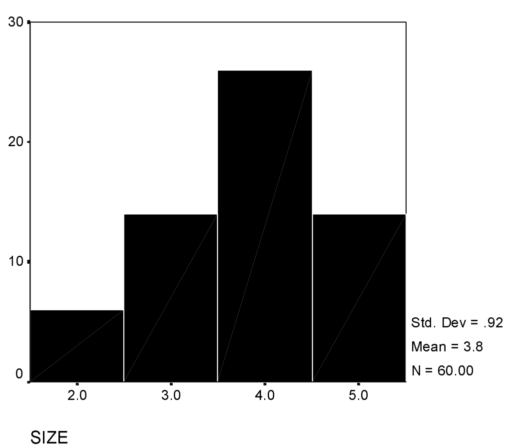

Figure E.3. Histogram of Size measure in Study 2.

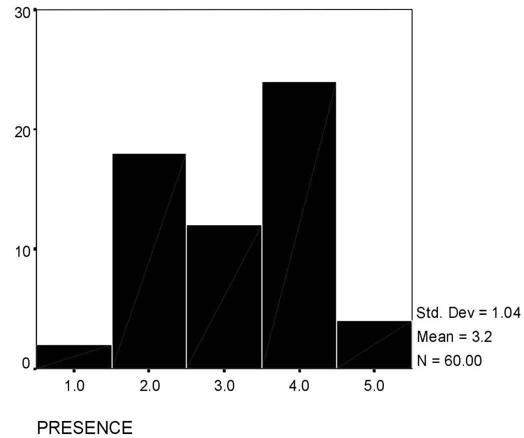

Figure E.4. Histogram of Presence measure in Study 2.

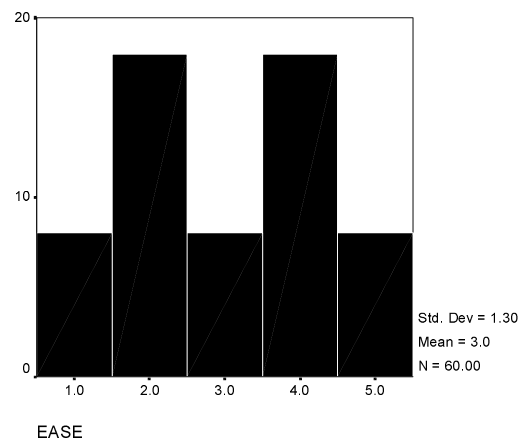

Figure E.5. Histogram of Ease measure in Study 2.

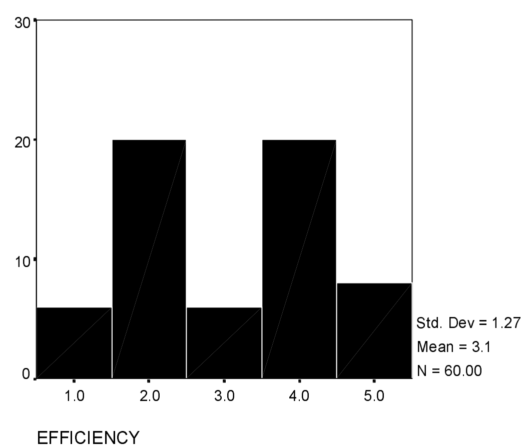

Figure E.6. Histogram of Efficiency measure in Study 2.

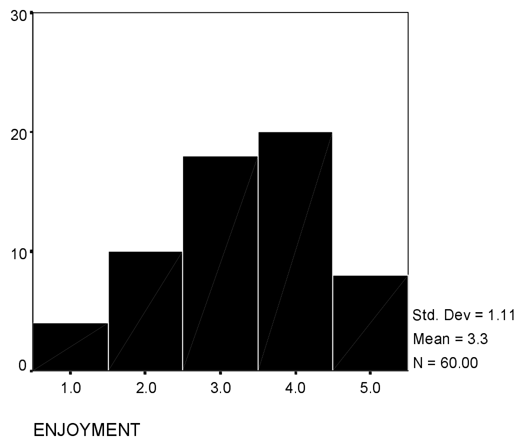

Figure E.7. Histogram of Enjoyment measure in Study 2.

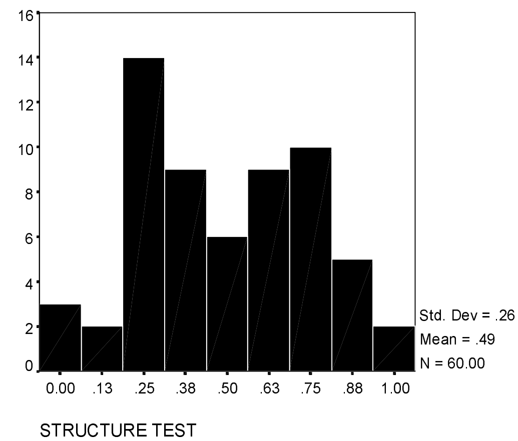

Figure E.8. Histogram of Structure Test measure in Study 2.

Appendix F

Study 3 – Design, Script, and Questionnaire

Subject	Block	Design	Targets
1	1	nat	1
1	2	efc	2
2	1	nat	2
2	2	efc	1
3	1	efc	1
3	2	nat	2
4	1	efc	2
4	2	nat	1
5	1	nat	1
5	2	efc	2
6	1	nat	2
6	2	efc	1
7	1	efc	1
7	2	nat	2
8	1	efc	2
8	2	nat	1

Table F.1. Experimental design for Study 3.

Parallel Worlds Project
Script for Study 3
Umeå University, July 1999

1) **Preparation**
   a) Restart computer
   b) blank notepad page
   c) Questionnaire with two parts and summary, labeled and ordered per subject
   d) Close blinds
   e) Close all unnecessary PC applications (e.g., email, ICQ)
2) **Introduction**
   a) Welcome
      i) Experimenter thanks subject for participation
      ii) Experimenter and subject introduce selves, if not acquainted already
   b) Explanation of experiment
      i) My Ph.D. research is about how people experience (desktop) virtual reality
      ii) For experiments, I developed some virtual worlds
      iii) Supporting hardware is a PC w/ good graphics card; software is CosmoPlayer VRML plug-in for Netscape
3) **Experiment instructions**
   a) Total time about 75 minutes
      i) 15 minutes VR training time
      ii) Exploring 2 virtual worlds + short questionnaire for each one
      iii) Short summary questionnaire at end
   b) In each virtual world
      i) First, 2 minutes of free exploration to begin
      ii) Then, hunt per list of targets to find as many as possible (in order)
         (1) Each list item shows a hint, plus the target's exact name in italics
            (a) There are more targets listed than anybody is expected to find
         (2) When a target is found
            (a) First, click crosshair icon on control panel (expect reminder)
            (b) Then, click on target's wall or balcony to play a sound
            (c) Finally, dbl-click crosshair icon on control panel (expect reminder)
         (3) Hints for finding targets (repeat every 3 minutes as necessary)
            (a) Survey the environment often by spinning around in a circle
               (i) Each circular region can be surveyed best from its center
            (b) All virtual buildings contain information, but you may need to approach them to see text
            (c) All targets present – hunt elsewhere if target not in current region
            (d) OK to skip target item(s) & move to next item in list.
               (i) Please tell experimenter.
      iii) Finally, a short questionnaire
   c) Questions are OK anytime, and breaks are allowed between worlds

4) **Computer Training and Practice**
    a) Experimenter explains and subject practices VR UI (10-15 minutes)
        i) Setup software
            (1) Start & maximize Netscape Communicator
            (2) Select a Helsinki model (first apartments, then free choise)
                (a) http://www.helsinkiarena2000.fi/demos.html
        ii) Basic navigation
            (1) Experimenter explains and demos
                (a) Best navigation technique is flying directly to goal
                    (i) mouse click (down, up) on object/text
                (b) Manual flying is sort of like driving car
                    (i) Fly forwards and backwards, and/or spin around
                    (ii) Depress mouse button + drag mouse (in VR window)
                    (iii) Feedback from "rubber band" line direction and size
            (2) Subject practices 2-3 minutes
        iii) Another navigation technique
            (1) Experimenter explains
                (a) Can move parallel to monitor screen
                (b) Fly up and down, and/or right and left
                (c) Depress *right* mouse button + drag mouse (in VR window)
            (2) Subject practices 2-3 minutes
        iv) Extra navigation features
            (1) Experimenter explains
                (a) Fast flying – depress shift key while dragging mouse
                (b) Return home – click button at bottom left of window
            (2) Subject practices 2-3 minutes
5) **Experimenter's tasks per world**
    a) Experimenter sets up Netscape
        i) Open appropriate virtual world
        ii) Double-click "continuous seek" icon on control panel
    b) Subject exposed to virtual world
        i) Subject explores freely two minutes, then presses "Home" button
        ii) Experimenter gives subject appropriate target list
        iii) Subject hunts for targets for 20 minutes
            (1) Experimenter warns subject after 18 minutes.
        iv) Experimenter notes targets found or skipped (and times) on list
    c) Experimenter gives subject appropriate questionnaire
        i) Subject receives and fills in 1st questionnaire page
        ii) Subject *then* receives and fills in 2nd questionnaire page
6) **Conclusion**
    a) Experimenter gives subject 50 crowns
    b) Experimenter requests subject's signature on payment form
    c) Experimenter thanks subject for participation in experiment

"Play" Targets 1

1. TV cartoon *Underdog*
2. The history of *Ancient Art*
3. Pop group *Spice Girls*
4. *Museums and Memorials* of military history
5. The Viking *Ericsson Leif 970 1020*
6. Baroque music *Ensembles*
7. Rock musician *Young Neil*
8. A Marvel comic book called *Thunderbolts*
9. Pop musician *Denver John 1943 1997*
10. *Violinists*
11. Internet Relay Chat (*IRC*) about music
12. The children's TV show *Clarissa Explains It All*
13. Philosopher *Foucault Michel 1926 1984*
14. TV police show *Homicide*
15. Musicians in *Luxembourg*
16. *Women s History Month*
17. Humorous music *Music That Sucks*
18. TV sketch comedy *Roundhouse*
19. *Dixieland* jazz musicians
20. History of the *South Africa* region
21. *Dub* (Reggae) musicans
22. Rock/pop musician *Stefani Gwen*
23. Ancient Greek *Alexander the Great*
24. Musicians in the *Ska* genre
25. *Masters* of photography

"Play" Targets 2

1. The comic book *Tick The*
2. TV soap opera *As The World Turns*
3. History of the *Asia* region
4. The comic book *Spawn*
5. The TV cartoon *Pinky and the Brain*
6. TV action show *Viper*
7. Maritime (shipping) history of the *Titanic*
8. *Native American* musicians
9. Music *Recording Equipment*
10. Funny *Song Parodies*
11. *Maps* of the Medieval Crusades
12. The DC comic book *Justice Society of America*
13. TV dramas *Wonder Years The & Zorro*
14. *Reader Reviews* of music
15. The Marvel comic book *Spider Man*
16. The animated cartoon *Speed Racer*
17. The *Impressionism* movement in art
18. Philosopher *Heidegger Martin 1889 1976*
19. The literary genre of *Storytelling*
20. *Indices* of Christian musicians
21. TV police show *21 Jump Street*
22. Children's TV show *Blue s Clues*
23. The film *Gone With The Wind*
24. Recipes for *Soups*
25. *Organizations* in 19th-century American history

“Play” Questionnaire 1

Please classify the following items on as many levels as you can, according to the electronic world you have just seen. It’s OK to guess. Here is an example for *Sailing*:

Recreation –> Outdoors -> Sailing

1. Rock/pop musician *Martin Dean*

2. American *Historical Societies*

3. TV dramas *Cape The & Avonlea*

4. Folk musician *Buckley Jeff 1966 1997*

5. The *Bibliographical Society* on the history of books and printing

(continued on next page)

"Play" Questionnaire 1 (cont.)

1) Please rate the information categories according to size.

a) Arts ______
b) Entertainment ______
c) News & media ______

Very small	= 1
Small	= 2
Medium	= 3
Large	= 4
Very large	= 5

2) Please estimate the length of your session: ______ minutes

"Play" Questionnaire 2

Please classify the following items on as many levels as you can, according to the electronic world you have just seen. It's OK to guess. Here is an example for *Sailing*:

Recreation –> Outdoors -> Sailing

1. TV *Cooking Shows*

2. Philosopher *Baudrillard Jean*

3. The comic book *Shi*

4. *Weekly* food recipes

5. *Disc Jockeys* in the reggae genre

(continued on next page)

"Play" Questionnaire 2 (cont.)

1) Please rate the information categories according to size.

a) Arts ______
b) Entertainment ______
c) News & Media ______

Very small	= 1
Small	= 2
Medium	= 3
Large	= 4
Very large	= 5

2) Please estimate the length of your session: ______ minutes.

Summary Questionnaire
Parallel Worlds Study 3
Umeå University
July 1999

1) Please rate the worlds according to size.

a) First world ______
b) Second world ______

Very small	= 1
Small	= 2
Medium	= 3
Large	= 4
Very large	= 5

2) Please rate the worlds according to sense of presence – “You are there.”

a) First world ______
b) Second world ______

Very weak	= 1
Weak	= 2
Medium	= 3
Strong	= 4
Very strong	= 5

3) Please rate the worlds according to ease of use.

a) First world ______
b) Second world ______

Very hard	= 1
Hard	= 2
Medium	= 3
Easy	= 4
Very easy	= 5

4) Please rate the worlds according to efficiency.

a) First world ______
b) Second world ______

Very inefficient	= 1
Inefficient	= 2
Neutral	= 3
Efficient	= 4
Very efficient	= 5

5) Please rate the worlds according to enjoyment.

a) First world ______
b) Second world ______

Very unpleasant	= 1
Unpleasant	= 2
Neutral	= 3
Pleasant	= 4
Very pleasant	= 5

Appendix G

Study 3 – Data Summary

		SUBJ	BLOCK	DES	DURA-TION	SIZE	TAR-GETS	PRES-ENCE	EASE	EFFIC	ENJOY	STRUCT TEST
1		1	1	nat	30.00	4.0	4	2	1	2	2.0	.40
2		2	1	nat	10.00	3.0	5	4	2	1	3.0	.70
3		5	1	nat	15.00	3.0	2	2	2	2	3.0	.60
4		6	1	nat	15.00	4.0	2	1	1	1	1.0	.50
5		3	1	efc	20.00	2.0	4	3	4	4	4.0	.50
6		4	1	efc	20.00	3.0	13	4	4	4	3.0	.80
7		7	1	efc	15.00	3.0	16	2	4	3	4.0	.80
8		8	1	efc	20.00	3.5	3	3	2	3	4.0	.60
9		3	2	nat	20.00	2.0	11	2	3	2	4.0	.80
10		4	2	nat	30.00	5.0	13	3	2	3	4.0	.60
11		7	2	nat	10.00	3.0	16	2	3	2	2.0	.80
12		8	2	nat	30.00	3.5	4	4	3	4	4.5	.60
13		1	2	efc	30.00	3.0	13	4	4	5	4.0	.70
14		2	2	efc	13.00	3.0	13	3	4	4	4.0	.80
15		5	2	efc	10.00	4.0	6	2	5	4	4.0	.80
16		6	2	efc	10.00	3.0	4	3	3	3	4.0	.60
Total	N	16	16	16	16	16	16	16	16	16	16	16

Table G.1. Case summaries for Study 3.

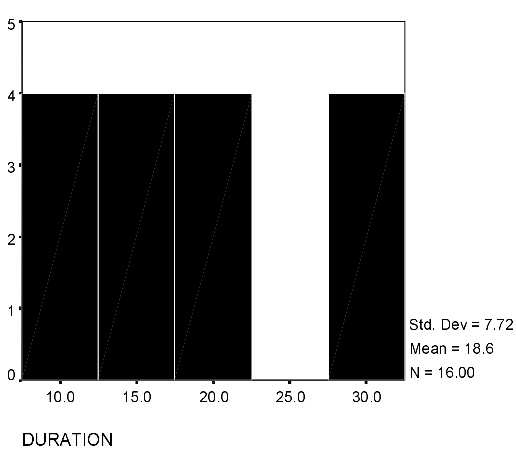

Figure G.1. Histogram of Duration measure in Study 3.

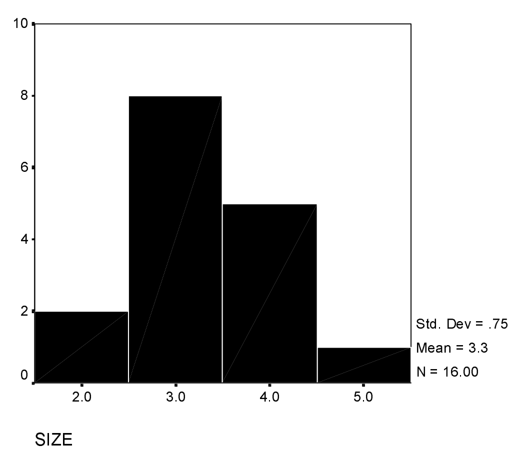

Figure G.2. Histogram of Size measure in Study 3.

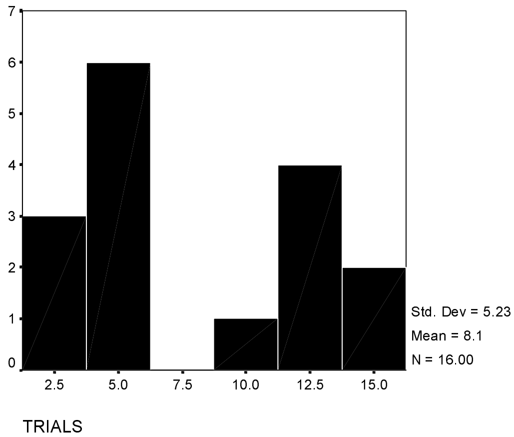

Figure G.3. Histogram of Targets measure in Study 3.

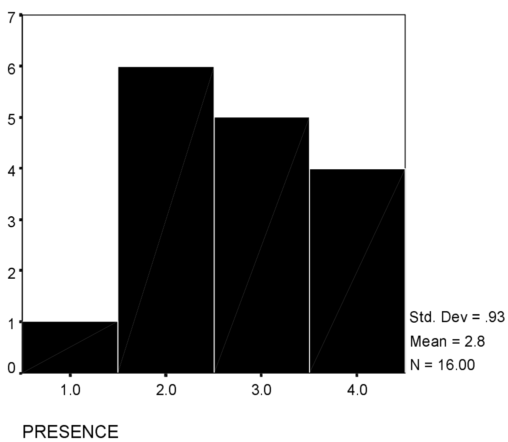

Figure G.4. Histogram of Presence measure in Study 3.

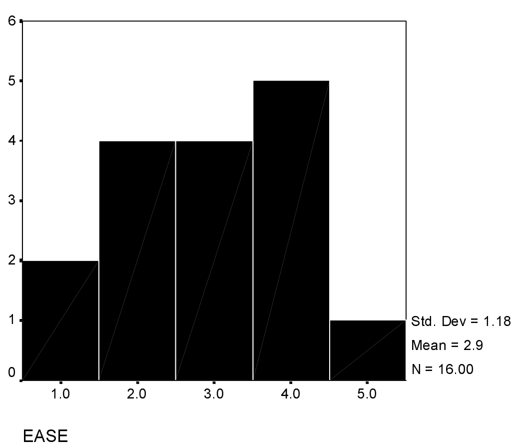

Figure G.5. Histogram of Ease measure in Study 3.

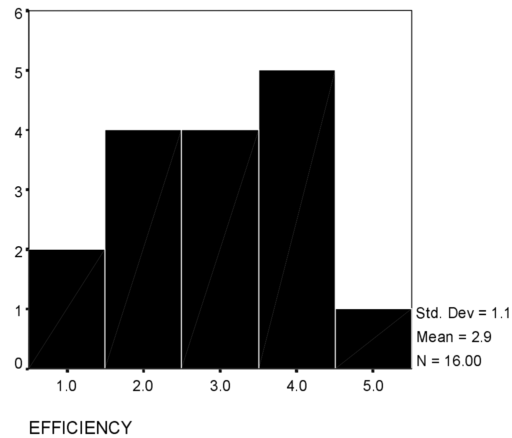

Figure G.6. Histogram of Efficiency measure in Study 3.

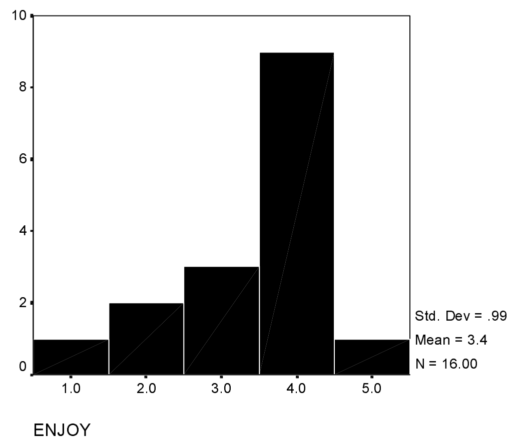

Figure G.7. Histogram of Enjoyment measure in Study 3.

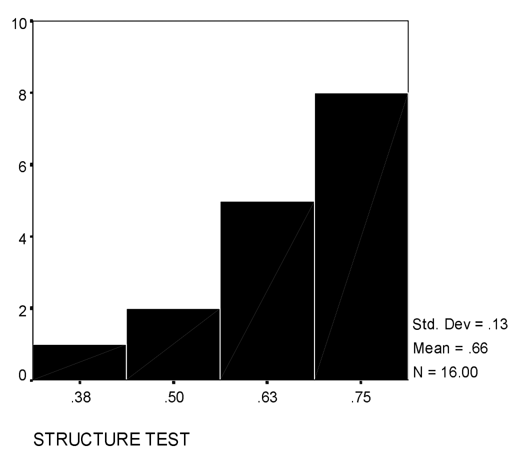

Figure G.8. Histogram of Structure Test measure in Study 3.

Appendix H

Study 4 – Design, Script, and Questionnaires

Subject	Block	Design
1	1	Map
1	2	Vr
2	1	Vr
2	2	Map
3	1	Map
3	2	Vr
4	1	Vr
4	2	Map
5	1	Map
5	2	Vr
6	1	Vr
6	2	Map
7	1	Map
7	2	Vr
8	1	Vr
8	2	Map
9	1	Map
9	2	vr
10	1	vr
10	2	map
11	1	map
11	2	vr
12	1	vr
12	2	map
13	1	map
13	2	vr
14	1	vr
14	2	map
15	1	map
15	2	vr
16	1	vr
16	2	map
17	1	Map
17	2	Vr
18	1	Vr
18	2	Map
19	1	Map
19	2	Vr
20	1	Vr
20	2	Map

Table H.1 Experimental design for Study 4.

David Modjeska
Dept. of MIE
U of T

Script for Study 4
30 January 2000
(approx. 2 hours)

1. Preparation
    1.1 Reboot PC; close all apps; open Netscape maximized w/o taskbars; close Netscape
    1.2 Close window blinds and pull out desk trays about an inch
    1.3 Assemble payment form, $20, notebook, stopwatch, sample target card, blank forms packet, movement controls sheet, and 2 pens
    1.4 Select two sets of 10 targets for this subject, per master list

2 Introduction and training (28 minutes)
    2.0.1 [Small talk to verify English-language fluency – e.g., weather, subject's work/study]
    2.1 General introduction – welcome and identification (2 minutes)
        2.1.1 Experimenter – David Modjeska, Ph.D. student, DCS
        2.1.2 Lab – Interactive Media Lab of Prof. Mark Chignell, MIE
        2.1.3 Experiment overview – intro, 2 online scavenger hunts w/ questionnaires (and break between hunts), and conclusion w/ questionnaire
        2.1.4 [If subject has not already been asked about Yahoo Web index, ask now]
            2.1.4.1 "For general background information, how often you use the Yahoo Web index?"
            2.1.4.2 Only if daily use - thank subject, pay $10, get signature, and dismiss them.
        2.1.5 Subject signs explanation and consent form (1.5 minutes)
        2.1.6 "As I proceed, questions are OK any time."

    2.2 Demo and practice in CosmoPlayer (22 minutes)

        2.2.1 Introduction (2 minutes)
            2.2.1.1 Project – info. visualization & user navigation in desktop VR (virtual reality)
            2.2.1.2 Training will use same worlds as scavenger hunts
                2.2.1.2.1 Each world is 3D model w/ water, islands, landmarks, and buildings
            2.2.1.3 World represents a Web index, without actual links to pages

        2.2.2 Fly though of sample world play_fly.wrl (12 minutes)
            2.2.2.1 Home viewpoint (double-click "continuous seek" icon)
                2.2.2.1.1 Can return to home viewpoint anytime by clicking "Home" button
            2.2.2.2 Go in model space
                2.2.2.2.1 Go forward and backward, by dragging vert. line with left mouse button
                2.2.2.2.2 Turn right and left, by dragging hz. line with left mouse button
                2.2.2.2.3 Big lines = fast movement, little lines = slow movement
                2.2.2.2.4 Accelerate going with shift key for long distances
                2.2.2.2.5 Note that object labels will pivot out of your way as they're encountered
                2.2.2.2.6 User practices navigation (up to 4 minutes)
            2.2.2.3 Slide in screen plane
                2.2.2.3.1 Slide up and down, by dragging vert. line with right mouse button
                2.2.2.3.2 Slide right and left, by dragging hz. line with right mouse button
                2.2.2.3.3 Accelerate sliding with shift key for long distances
                2.2.2.3.4 User practices navigation (up to 4 minutes)
            2.2.2.4 Click-and-jump
                2.2.2.4.1 Single-click on any object or label to jump directly to it
                2.2.2.4.2 User practices navigation (up to 4 minutes)
            2.2.2.5 Close Netscape (and then count to 10)

2.2.3 Map view of sample world play_map.wrl (8 minutes)
2.2.3.1 (Double-click "continuous seek" icon)
2.2.3.1.1 Like paper map plus camera zoom
2.2.3.1.2 Navigate like in fly-through world
2.2.3.1.2.1 But before go or jump to obj or label, helps to center it on screen first
2.2.3.1.3 If map becomes tilted, just turn right or left to straighten it
2.2.3.1.4 User practices (up to 8 minutes)
2.2.3.1.5 Close Netscape (and then count to 10)

2.3 Target-hunt introduction (4 minutes)

2.3.1 Procedure (2 minutes)
2.3.1.1 There will be two hunts, each lasting 20 minutes
2.3.1.2 Subject will hunt for up to 4-10 targets, depending on subject's speed
2.3.1.3 For each target, subject will receive card (show sample) & up to 5 mins.
2.3.1.3.1 When card is rec', *subject presses "OK" & Home* buttons on screen
2.3.1.3.2 N.b., each card shows a general hint or general context
2.3.1.3.3 Plus specific target topic in bold, italic text (before or after card's hint)
2.3.1.3.4 A topic can appear in several places in world, but only 1 context fits card
2.3.1.4 In fly-through world, target is bldg. floor; in map-view, target is on bldg. roof
2.3.1.5 When subject finds target, show it to experimenter for confirmation.
2.3.1.6 After each hunt, subject will fill in classification form for 5 hypothetical targets
2.3.1.6.1 Model world will be hidden
2.3.1.6.2 Show sample classification form

2.3.2 Exploration hints (2 minutes)
2.3.2.1 Each object has a label, sometimes visible only near object
2.3.2.2 There are no hidden objects in the world
2.3.2.3 Overview helpful - best to survey a region by spinning or moving at its center
2.3.2.4 Items in each circle are in alphabetical order
2.3.2.5 All targets in world – if not where expected, look elsewhere in world
2.3.2.5.1 Category "Humanities" includes history & literature, among other things
2.3.2.5.2 Category "Media" includes television & radio, among other things

3 Target Hunt # 1 (30 minutes)
3.0 If subject number is odd, use App_Map.html, else App_Fly.html
3.0.1 (Double-click "continuous seek" icon) (Show "OK" button)
3.1 Free exploration (1 minute)
3.2 Hunt in "play" world (20 minutes)
3.2.1 4-10 targets @ 2-5 minutes/target
3.2.2 Experimenter records each target outcome and hunt time in notebook, using stop-watch
3.2.3 Close Netscape (and then count to 10)
3.3 Questionnaire # 1 (10 minutes)
3.3.1 Minimize browser window during questionnaire
3.3.2 "You will have ten minutes to complete this questionnaire. But first, take a minute to read the instructions, and ask any questions. The questionnaire is about classifying 5 hypothetical targets from general category down to specific topic."

4 Break (10 minutes)
4.0 "Let us take a break of 5 to 10 minutes."
4.1 Rename vrmlLog.txt to subj[subject #]a.txt & move to Navig_Log folder

5 Target Hunt # 2 (30 minutes)

- 5.0 If subject number is odd, use App_Fly.html, else App_Map.html
  - 5.0.1 (Double-click “continuous seek” icon)
  - 5.0.2 ”The second hunt will use the same procedures as the first hunt.”
- 5.1 Free exploration (1 minute)
- 5.2 Hunt in “play” world (20 minutes)
  - 5.2.1 4-10 targets @ 2-5 minutes/target
  - 5.2.2 Experimenter records each target outcome and hunt time in notebook, using stop-watch
  - 5.2.3 Close Netscape (and then count to 10)
- 5.3 Questionnaire # 2 (10 minutes)
  - 5.3.1 Minimize browser window during questionnaire
  - 5.3.2 “You will have ten minutes to complete this questionnaire. But first, take a minute to read the instructions, and ask any questions. The questionnaire is about classifying 5 hypothetical targets from general category down to specific topic.”

6 Conclusion (10 minutes)

- 6.0 Questionnaire # 3 (5 minutes)
- 6.1 Subject signs payment and receives $20
- 6.2 Experimenter debriefs subject and answers any questions
  - 6.2.1 “I are looking at how to create a visual representation of information structure, so that people can find things easily and also learn where things are. In this experiment I used a part of the Yahoo! hierarchy as the information, and then I asked you to find targets under 2 different conditions. One condition allowed you to fly through the world in 3 dimensions, while the other condition gave you a map view of the world with zooming. Based on the results that I get, I hope to make recommendations about how designers should build virtual worlds, where people are looking for particular information in a large structure.”
- 6.3 Experimenter thanks subject

7 Clean-up

- 7.0 Rename vrmlLog.txt to subj[subject #]b.txt & move to Navig_Log folder
- 7.1 Replace 20 target cards in stack
- 7.2 Open window blinds

Questionnaire # 1

Please classify the 5 sample items below (A – E) as accurately as you can, according to the virtual world you have explored. Guessing is permitted.

When classifying each item:

(1) Choose 6 words or phrases. The 2 long blank lines are available for notes as needed.
(2) Write the same 6 words or phrases, in order from general category to specific topic, on the 6 short blank lines

For example, here is a classification of the item "Internet Radio ***Shows*** ."

1. Computers & Internet
    2. Internet
        3. Internet Broadcasting
            4. Radio
                5. Stations and Shows
                    6. ***Shows***

A. History of ***Hungary***

____

____

1. __________________________
2. __________________________
3. __________________________
4. __________________________
5. __________________________
6. ***Hungary***

B. Musicians from ***Croatia***

____

____

1. __________________________
2. __________________________
3. __________________________
4. __________________________
5. __________________________
6. ***Croatia***

C. History ***Magazines*** about Geneology

____

____

1. ______________________
2. ______________________
3. ______________________
4. ______________________
5. ______________________
6. ***<u>Magazines</u>***

D. ***Gregorian Chant*** Music of the Middle Ages

____

____

1. ______________________
2. ______________________
3. ______________________
4. ______________________
5. ______________________
6. ***<u>Gregorian Chant</u>***

E. DC Comic Book ***Sandman*** (on the Vertigo label)

____

____

1. ________________________

2. ________________________

3. ________________________

4. ________________________

5. ________________________

6. ***Sandman***

Questionnaire # 2

Please classify the 5 sample items below (A – E) as accurately as you can, according to the virtual world you have explored. Guessing is permitted.

When classifying each item:

(3) Choose 6 words or phrases. The 2 long blank lines are available for notes as needed.
(4) Write the same 6 words or phrases, in order from general category to specific topic, on the 6 short blank lines

For example, here is a classification of the item "Internet Radio ***Shows*** ."

1. Computers & Internet
    2. Internet
        3. Internet Broadcasting
            4. Radio
                5. Stations and Shows
                    6. ***Shows***

A. TV Shows about ***Literature***

____

____

1. ______________________

2. ______________________

3. ______________________

4. ______________________

5. ______________________

6. ***Literature***

B. New Age Musician ***Vangelis***

____

____

1. ______________________

2. ______________________

3. ______________________

4. ______________________

5. ______________________

6. ***Vangelis***

C. ***Pre-Raphaelites*** Period in Art History

____

____

1. ____________________________
2. ____________________________
3. ____________________________
4. ____________________________
5. ____________________________
6. ***Pre-Raphaelites***

D. ***Sewing*** as a Design Art

____

____

1. ____________________________
2. ____________________________
3. ____________________________
4. ____________________________
5. ____________________________
6. ***Sewing***

E. TV Comedy ***The Naked Truth***

____

____

1. ____________________________
2. ____________________________
3. ____________________________
4. ____________________________
5. ____________________________
6. ***The Naked Truth***

Questionnaire # 3

For each question below, please circle the best response (e.g., "Medium")

1) Please rate each world according to sense of presence – "You are there."

Map-view:	*Fly-through:*
Very weak Weak Medium Strong Very strong	Very weak Weak Medium Strong Very strong

2) Please rate each world according to ease of use.

Map-view:	*Fly-through:*
Very hard Hard Medium Easy Very easy	Very hard Hard Medium Easy Very easy

3) Please rate each world according to efficiency.

Map-view:	*Fly-through:*
Very inefficient Inefficient Medium Efficient Very efficient	Very inefficient Inefficient Medium Efficient Very efficient

4) Please rate each world according to enjoyment.

Map-view:	*Fly-through:*
Very unpleasant Unpleasant Neutral Pleasant Very pleasant	Very unpleasant Unpleasant Neutral Pleasant Very pleasant

Appendix I

Study 4 – Data Summary

		SUBJ	BLOK	DES	MINN	PRESENCE	EASE	EFC	JOY	TRIALS	TRAVEL	PROX	ERRS	STRUCT TEST
1		1	1	1	46	4	1	4	3	7	3.38	1.43	2.71	8
2		1	2	2	46	3	2	4	3	8	3.02	1.13	2.88	18
3		2	1	2	37	4	4	4	5	6	2.87	1.45	2.17	11
4		2	2	1	37	2	2	3	3	6	4.42	1.17	3.50	18
5		3	1	1	57	4	3	4	3	7	1.85	1.77	1.14	21
6		3	2	2	57	3	4	3	5	13	1.51	1.36	1.54	15
7		4	1	2	38	2	3	3	4	5	4.13	1.13	3.00	23
8		4	2	1	38	3	4	4	3	8	2.23	1.83	1.38	22
9		5	1	1	41	4	3	3	3	11	1.75	2.22	2.27	20
10		5	2	2	41	2	2	3	1	12	2.67	1.69	1.33	25
11		6	1	2	37	4	2	4	5	6	2.42	1.64	4.83	14
12		6	2	1	37	3	4	4	4	11	1.53	1.76	.64	21
13		8	1	2	52	5	2	2	4	7	4.29	1.49	5.43	18
14		8	2	1	52	2	4	5	3	15	1.53	2.11	1.27	27
15		9	1	1	48	3	2	3	4	19	1.24	2.28	1.32	33
16		9	2	2	48	4	3	4	5	23	1.19	2.03	.57	33
17		10	1	2	50	5	2	3	4	7	2.41	1.30	1.43	22
18		10	2	1	50	3	4	4	3	14	1.58	2.19	.86	27
19		11	1	1	39	3	3	3	3	15	1.01	2.28	1.00	20
20		11	2	2	39	4	3	3	4	17	1.61	1.80	1.12	21
21		12	1	2	48	4	2	4	4	18	2.17	1.96	1.28	28
22		12	2	1	48	3	2	4	2	26	.87	2.31	.62	30
23		13	1	1	31	3	3	5	4	7	1.89	2.21	2.14	12
24		13	2	2	31	5	4	3	5	6	6.53	1.55	7.33	15
25		14	1	2	51	4	2	2	4	15	1.91	1.58	1.20	24
26		14	2	1	51	3	3	4	3	29	.95	2.46	.86	29
27		15	1	1	44	4	3	3	3	13	1.48	2.23	1.54	18
28		15	2	2	44	1	1	1	1	11	1.86	2.11	3.18	22
29		16	1	2	41	5	4	3	4	13	2.21	1.82	2.62	22
30		16	2	1	41	3	2	5	5	18	1.34	2.13	.72	31
31		17	1	1	23	4	2	2	2	8	3.39	1.65	4.50	32
32		17	2	2	23	3	4	4	4	12	2.72	1.79	3.58	23
33		18	1	2	40	4	5	1	4	12	1.86	1.67	1.17	23
34		18	2	1	40	2	3	2	2	16	1.86	2.02	.94	23
35		19	1	1	49	4	4	5	4	10	3.06	1.66	1.60	24
36		19	2	2	49	2	2	1	1	7	3.60	1.33	3.57	20
37		20	1	2	49	3	4	3	4	7	2.78	1.27	.86	27
38		20	2	1	49	4	4	4	4	10	1.49	1.94	.70	23
39		21	1	1	58	4	4	4	4	22	1.34	2.18	.45	20
40		21	2	2	58	5	3	3	4	14	2.39	1.59	1.36	21
Total	N	40	40	40	40	40	40	40	40	40	40	40	40	40

Table I.1. Study 4 – Case summaries for Study 4.

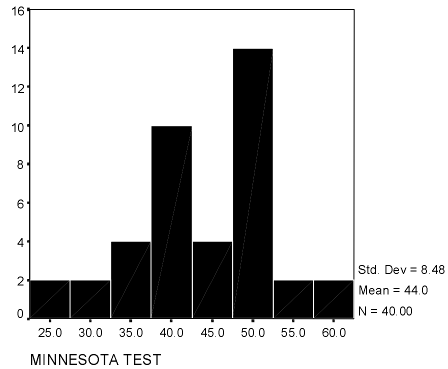

Figure I.1. Histogram of Minnesota Test measure in Study 4.

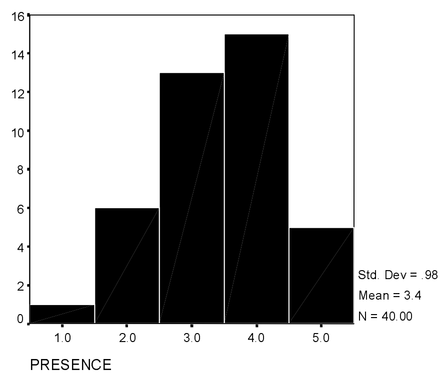

Figure I.2. Histogram of Presence measure in Study 4.

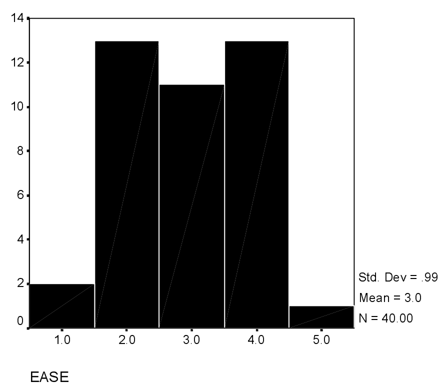

Figure I.3. Histogram of Ease measure in Study 4.

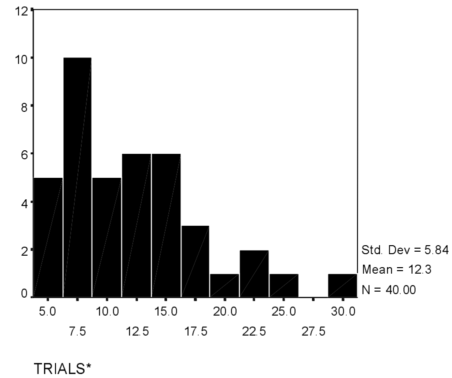

Figure I.4. Histogram of Trials measure in Study 4.

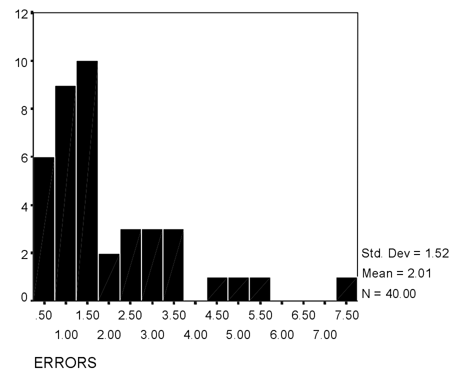

Figure I.5. Histogram of Errors measure in Study 4.

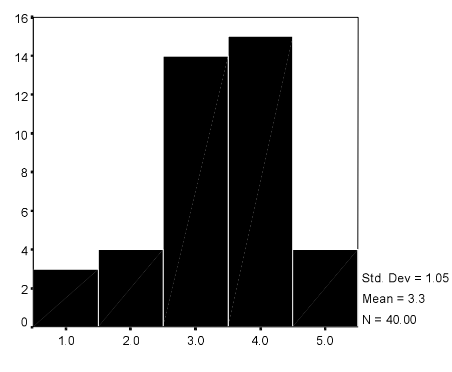

Figure I.6. Histogram of Efficiency measure in Study 4.

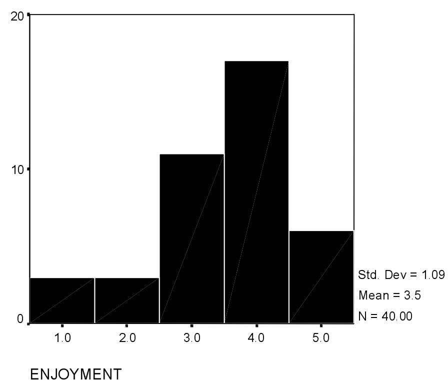

Figure I.7. Histogram of Enjoyment measure in Study 4.

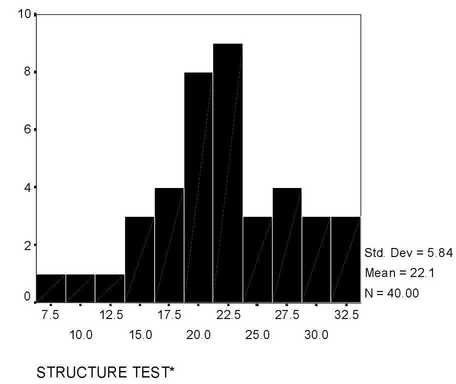

Figure I.8. Histogram of Structure Test* measure in Study 4.