Navigation in Electronic Worlds:
Research Review for Depth Oral Exam

David Modjeska
Department of Computer Science
1 May 1997

1. INTRODUCTION

Navigate: **1.a.** To go from one place to another in a ship or ships, to sail. **. . . c.** To walk steadily; to keep on one's course. . . . **6.a.** To manage, direct, sail, or fly . . . **b.** To travel or fly through . . . .
(from the *Oxford English Dictionary - 2nd Edition*)

Navigation: the process of determining a path to be traveled by any object through any environment (from an article by Darken and Silbert on wayfinding in virtual worlds [21])

1.1 General

Navigation is a fundamental human activity. It occurs in a wide range of activities and contexts, and it has accordingly been studied by a range of disciplines. In the social sciences, anthropologists have studied the tools and techniques of different cultures, including Western ones. Psychologists have studied the cognitive structures and processes involved in human navigation. In the field of urban design, architects and urban planners have studied the relationship between the built environment and navigational activities, particularly with an eye to improving both. In the applied and information sciences, human-computer interaction (HCI) and related researchers continue to investigate human information-seeking in large-scale information environments. Such environments tend to be either document-based hypermedia such as the World-Wide Web (WWW or Web) or continuous, spatial worlds such as virtual reality (VR). Both approaches are widespread in industrial applications and academic research.

Hypermedia models for information worlds, on the one hand, are discrete. (Hypermedia generalizes of hypertext to multiple media.) They take advantage of well-developed human skills in organizing and classifying information. Conveniently, hypermedia structures match basic ones in computer science, which facilitates implementation. For users, hypermedia offer familiar textual metaphors, as well as naturally supporting semantic realism. Hypermedia offer flexible, extensible, and powerful ways to organize information. Unfortunately, this approach sometimes has problems representing physical relationships.

(External) spatial models for information worlds, on the other hand, are continuous. They take advantage of well-developed human cognitive and neuro-motor skills for dealing with the physical world. Usefully, much research from the physical world can be adapted for electronic ones. For the user, spatial models offer intuitive physical metaphors, as well as naturally supporting physical realism. As with hypermedia, spatial models offer flexible, extensible, and powerful ways to organize information. Unfortunately, this approach sometimes has difficulty representing domain semantics.

As noted by numerous popular and academic commentators, information environments play an increasingly important role in our daily lives. These environments organize large amounts of information for both business and leisure activities. Hypermedia environments, on the one hand, include the rapidly-expanding Internet and corporate intranets; spatial environments, on the other hand, include numerous three-dimensional computer-graphics games. As these environments increase in scale and complexity, navigation becomes a significant issue for information users. Research and anecdotal evidence both note the frequency of user disorientation, which results in losing the sense of one's information-viewing history, current context, and/or ability to achieve information goals. In hypermedia, one can become "lost in hyperspace"; in spatial worlds, one can simply get lost. Given the importance of navigation in electronic worlds, the HCI community has recently increased its interest in this area. A workshop at the CHI '97 conference, for example, gathered researchers to consider navigation in electronic worlds. Major sub-issues include user tasks, information environments, navigational tools, user strategies, and individual differences between users. Another issue is the adaptation of social-scientific and urban-design research to electronic worlds.

A poster by the previously-mentioned CHI workshop offered a useful definition of navigation in electronic worlds: locomotion + wayfinding. Locomotion is the activity of moving from one location to another. Wayfinding is primarily a cognitive process, comprising three abilities [61]:

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

That is, wayfinding is a spatial problem-solving process. We could define "skilled wayfinding behavior to be purposeful, oriented movement during navigation" [22]. Although it will consider navigation in general, this research review will focus on wayfinding, concentrating on the aspect of navigation that applies to the broadest range of information environments.

1.2 *Organization of Review*

For organizing research related to navigation in electronic worlds, several approaches are possible. One has been chosen for organizational appropriateness and narrative convenience. The remaining sections of the paper will cover background issues, physical and information exploration, information structuring, 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. The section on physical and information exploration will introduce key research on navigation. It describes general models, introduces the concept of mental maps, discusses wayfinding research, presents an anthropological perspective, and reviews some theory and experiments in hypermedia. The section on information structuring contributes to the paper's topic in two ways. First, it discusses the properties and benefits of specific information structures, which affect both locomotion and wayfinding in electronic worlds. Second, it suggests structural principles that could facilitate online navigation. The section on information visualization contributes in two ways to the research paper. It discusses a variety of techniques and tools that could facilitate electronic navigation. It then considers 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 general conclusions section, finally, will summarize key research issues and directions for future research.

2. PERCEPTION AND COGNITION

Before discussing the main topics of this research review, it is worth reviewing briefly some key 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.

2.1 Cognitive Engineering

Norman notes that cognitive engineering is not cognitive psychology, cognitive science, or human factors (ergonomics). Rather, "[I]t is a type of applied cognitive science, trying to apply what is known from science to the design and construction of machines" [58]. Cognitive engineering thus has two major goals:

1. understanding the principles of human action and performance that are relevant for developing engineering design principles
2. designing systems that are pleasant to use, as well as efficient, easy, and powerful.

A central concept of cognitive engineering, mental models are described as follows:

> Mental models seem a pervasive property of humans. . . . people form internal, mental models of themselves and of the things and people with whom they interact. These models provide predictive and explanatory power for understanding the interaction. Mental models evolve naturally through interaction with the world and with the particular system under consideration. [58].

In discussing mental models, Normal distinguishes between four related concepts [59]:

- the target information system
- a developer's conceptual model of the target system
- the user's mental model of the target system
- the scientists conceptualization of the mental model

Six general observations can be made about mental models: they are incomplete; human abilities to "run" them are limited; they are unstable; they have no firm boundaries; they are "unscientific" or "superstitious"; and they are parsimonious [59].

The dominant paradigm for UI design, as reflected in the this review, is that of direct manipulation. Shneiderman proposes the following definition [70]:

- Continuous representation of the objects and actions of interest,
- Physical actions or presses of labeled buttons instead of complex syntax,
- Rapid incremental reversible operations whose effect on the object of interest is immediately visible.

According to Hutchins *et al.* [39], 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. (According to our earlier definition of navigation, locomotion is an activity at the articulator level, while wayfinding as an activity at the semantic level.) 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" [39]. In summary, the most direct UI will minimize distance while maximizing engagement.

2.2 *UI METAPHORS*

Although several researchers favor HCI interactions "[l]ike interacting with objects in the physical world" [37], 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 claim that, "our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature" [43]. They claim that "the essence of metaphor is understanding and experiencing one kind of thing in terms of another" [43]. 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 are believed to 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.

Through the power of similarity, UI metaphors can help users to understand and predict the behavior of an information system. A metaphor is a type of mental model. Examples of common UI metaphors include file folders, trash cans, and paint palettes, as well as sheets of paper for calculating, writing, and drawing. Mountford notes that metaphors facilitate creativity in UI designers [53]. She suggests several techniques for generating new metaphors: re-conceptualization, adaptation, modification, combination, and/or replacement of UI objects. Erickson offers further suggestions for designing UI metaphors [26]. He notes the need for a functional definition of the target system, as well as an identification of users' problems. For evaluating the utility of metaphors, he suggests five criteria: amount of structure, applicability of structure, representability, suitability to audience, and extensibility.

3. PHYSICAL AND INFORMATION EXPLORATION

This section will review research on navigation and information-seeking behavior. 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, mental maps, wayfinding, anthropological perspective, and hypermedia environments.

3.1 *GENERAL MODELS*

For information systems, Waterworth and Chignell propose a general model of user-initiated exploration [82]. 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 model is called querying; if the user has a less specific target and some interest in serendipity, the model 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. While useful, the proposed model could be improved by taking greater cognizance of the user's understanding of the information system [33].

Pirolli and Card have developed a model of the cost structure of information navigation and gathering [64]. Information foraging theory is used to analyze human activities with information access technologies. The theory is adapted from optimal foraging theory in biology and anthropology, which 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" [63].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" [64]. 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 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.

Information foraging theory has three implications for navigation in electronic worlds. First, it suggests that information cost structure should be considered explicitly in the design of online environments. Second, an electronic worlds (or sub-world) could 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 (especially [61]).

3.2 MENTAL MAPS

Cognitive maps are essential for human storage and use of environmental information, e.g., during navigation. Such maps are the dominant concept for internal representations of large-scale environments. The term "cognitive map" was proposed by Tolman [78]. To counter behaviorist theories of stimulus-response bases for animal navigation, Tolman conducted a series of experiments with rats in mazes. His research demonstrated that rats navigated not just using past stimuli, but according to learned 2D spatial position. Tolman proposed that "in the course of learning something like a field map of the environment gets established in the rat's brain" [78]. Tolman described this knowledge as a "cognitive-like map of the environment . . . indicating routes and paths and environmental relationships" [78]. He speculated that such vary between narrow strips and broader, comprehensive views. Later research by Thorndyke and Hayes-Roth, and others, supported this speculation [77]. Researchers continue to debate whether the representations of cognitive maps are analog (continuous) or propositional (discrete). Evidence for propositional representations is strong, however, as discussed below.

Research by Stevens & Coupe provided evidence of hierarchical representations of environmental knowledge [73]. 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 to 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 [19]. 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. 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 concept of mental maps was extended by Downs and Stea to encompass definition, development, and use [25]. Their compact 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 . . . [25]

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." 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.

Key experimental research by Thorndyke and Hayes-Roth shows differences in spatial knowledge acquired from maps and physical navigation [77]. The authors propose a broad distinction between two types of spatial knowledge, which we will preface with a third type mentioned in [50]:

- basic feature recognition or *landmark knowledge*
- 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 between a physical map and its mental representation. This assumption allows them to compare internal and external maps more effectively; the assumption is not contradicted in their study. 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 conducted experiments with forty-eight subjects in the large headquarters of the Rand Corporation [75]. 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, on account 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 environmental 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 the user errors described by Stevens and Coupe [73], and by Chase [19].

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 consider two key questions [13]. 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" [13]. 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 difference in individual spatial ability are shown by 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.

3.3 WAYFINDING

3.3.1 PHYSICAL WORLD

In research that has influenced the fields of urban design, psychology, and computer science (information visualization), urban planner Kevin Lynch discovered psychologically important, structural patterns that improve complex information environments [45]. 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." [45]

To investigate the images held by city residents, Lynch conducted research in three American cities. This research involved resident interviews, map sketches (cognitive cartography), trip and feature `descriptions, and photo classification; trained observers also conducted field analysis. These disparate data were assembled into cartographic maps. Lynch's methodology is notable for assembling heterogeneous research signatures (sets of encoding and decoding operations per [25]) into unified cartographic maps. Lynch's 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.
- *Edges*, other linear elements, act as boundaries between two areas.
- *Districts* are two-dimensional city sections into which observers mentally enter, and which have distinctive characters.
- *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.
- *Landmarks* are single objects external to viewers, and upon which longtime city residents rely strongly.

Lynch notes that a satisfying city form should weave these five image elements into a strong pattern. The resulting overall image usually comprises a set of smaller images which 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 ten formal qualities upon which a designer may operate to create image elements of the recommended types. These guidelines may be found in [45]. 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 topological in nature, Lynch's structural recommendations for the physical world may improve wayfinding in virtual worlds as well. Research is needed to validate this claim (e.g., [17], [40]). In this vein, Shum notes that users' cognitive maps of hypermedia clearly include landmarks, paths, and nodes; districts and edges could be usefully added via UI metaphors [72]. Finally, Lynch's rich research methodology could be adapted and extended for electronic worlds (e.g., [22], [52]).

Later research by Passini built on Lynch's wayfinding research [61]. 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 (repeated from the introduction to this paper), 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

Wayfinding has both semantics (units and structures) and syntax (process). The semantics include tasks or goals, environmental information, plans or decisions, and actions. Environmental information can be descriptive, locational, and temporal; it involves perception and cognition, as well as cultural and personal differences. Plans have two general characteristics: a logical hierarchy of modular components, and ongoing temporal plans. These plans are developed iteratively and often use templates. 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 [77]. 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. Signs can offer three types of information: directional, identification, or confirmation; in general, signs can be verbal, pictorial, or symbolic.

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 environmental features helps people to develop 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:

1. cardinal or relative orientation of the map to the setting,
2. locating the destination

3. developing and memorizing a route.

For developing survey knowledge, the map should feature Lynch's five elements and Passini's three elements. In addition to the preceding ideas and observations, Passini offers specific guidelines for overall wayfinding design, architectural conception, signs, and maps; details can be found in [61].

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 recent research on information scent or residue (trace) by Pirolli and Card [63] and Furnas [29].

3.3.2 Virtual Worlds

Recent researchers have adapted physical wayfinding principles for virtual worlds, with mixed results. Two studies by Darken and Silbert support this adaptation ([21], [22]), but a study by Satalich argues against it [70].

Darken and Silbert investigated the problem of VR users' maintaining knowledge of locational information, that is, current position and orientation [21]. 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. In terms of size, a world is considered small when most or all of it can be seen from one viewpoint; a world is large when there is no viewpoint from which all of it can be seen. With regard to density, a world is considered sparse when it has a small number of objects and cues to assist navigation; it is dense when it has a large number of objects to assist navigation. In terms of activity, a world is considered static when the positions and values of objects do not change with time, thereby simplifying navigation; a world is complex when objects move (deterministically or randomly), thereby making navigation more complex. As part of their study, Darken and Silbert implemented a tool set of techniques based on physical navigation: flying, spatial audio, bread-crumb (history) markers, (grid) coordinate feedback, districting, landmarks, (Cartesian or polar) grid navigation, and map views (either forward-up or north-up).

Darken and Silbert conducted an informal experiment 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: (global) exploration, naïve search, and informed search. 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, subject 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. 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*. Details of the study are reported in [21]. 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 Silbert later conducted a more involved experiment that applied physical wayfinding and environmental design principles to virtual worlds [22]. 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:

1. divide the world into parts, preserving a sense of "place"
2. organize the parts under a simple principle
3. provide frequent directional cues.

Map principles, second, should present spatial information in a "flexible, orientation-independent representation of the environment" [22]:

1. show all organizational elements and principles,
2. always show the observer's position
3. orient the map forwards-up with respect to the observer.

As above, the authors classify wayfinding tasks into three categories: naïve (exhaustive) search, primed (non-exhaustive and informed) search, and exploration (without target).

Reflecting these principles, Darken and Silbert 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 virtual environments showing sea and land masses. Each treatment required five naïve 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 treatment, 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 applied to virtual worlds*. Moreover, human conceptions of virtual spaces may be analogous to conceptions of real spaces. Future research is needed in virtual environments with different spatial features.

Recent research by Satalich studies navigation and wayfinding in immersive VR, with the goal of designing tools and environmental cues to enhance navigational awareness. Several questions drive 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 follows the research of Thorndyke and Hayes-Roth [77]. To answer these questions, Satalich conducted a study with sixty-five subjects. The virtual environment (VE) represented a building 100 feet square, with a ceiling height of 10 feet and thirty-nine rooms. The building contained five-hundred objects to be used as landmarks. The experiment had a 3x2x2 between-subjects design. The first factor was the type of exploration, which lasted thirty 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 subject 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 subject'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.

Surprisingly, 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. Detailed experimental results available in [70]. According to the author, 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 Silbert 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 Silbert 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 Silbert 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 Silbert'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 Silbert'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.

Clarifying some issues of navigational tools and human psychology, Wickens reviews research in spatial perception and cognition and the display of spatial information [84]. Key topics include the acquisition of navigational information and spatial representations of non-spatial systems. As discussed above, people acquire navigational knowledge in three stages: landmark, route, and survey ([49], [77]). 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 proposes the principle of "visual momentum":

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

Wickens' comments apply to the research of Darken and Silbert [21] and Satalich [70]), as well to future work in this area. According to the principle of visual momentum, the VEs of Darken and Silbert, 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 Silbert 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.

3.4 Anthropological Perspective

From the global perspective of cognitive anthropology, Hutchins considers the tools and practices of team navigation ([37], [38]). Like Lynch and Passini, Hutchins conducted field research; unlike them, he used observational rather than experimental methods. For Hutchins, navigation answers the central question, "Where am I?". The answer establishes a correspondence between the surrounding world and some representation of the world. In any navigational system, questions are answered by combining one-dimensional constraints on position. The art of navigation involves the integration of information from multiple, simultaneous constraints to generate a solution. A map (or chart) is essentially a spatial analogy; it is an analog computational device. In different navigational traditions, different representational systems have different computational properties and allow different implementations.

For his research, Hutchins observed Micronesian navigators in canoes and American navigation teams on naval ships ([37], [38]). He compared their two sets of representational assumptions. In Western navigation, space is represented in miniature on a map. Earth coordinates provide the primary frame of reference. Direction, position, and distance are defined in a universal framework; combining universal time with distance gives universal speed. Directions, positions, distances, and speeds can be represented as numbers; all but speeds can be modeled in a chart. Navigational constraints (line-of-position, circles of position, and position-displacement) can be represented directly on a chart. Distance, rate, and time are represented as numbers, and constraint computations proceed by arithmetic. "All the major computations in this system are based on procedures that involve measurement (which is analog-to-digital conversion), followed by digital manipulation, followed by digital-to-analog conversion in the plotting of results on a chart" [37] In Micronesian navigation, by contrast, position is established by two lines of position: implicitly from a known location to the canoe, and imaginatively from the canoe to a conceptual island. Bird patterns and land sightings add circle-of-position constraints. The distance-rate-time constraint is represented in the superimposition of temporal landmarks on spatial landmarks, which are defined by the star bearings of the conceptual island. "In this system there are no universal units of direction, position, distance, or rate, no analog-to-digital conversions, and no digital computations. Instead, there are many special-purpose units and an elegant way of 'seeing' the world in which internal structure is superimposed on external structure to compose a computational image device" [37].

Hutchins notes that Western and Micronesian traditions diverged on account of three Western trends: (1) the increasing crystallization of knowledge and practice in the physical structure of artifacts, (2) the development of measure as an analog-to-digital conversion, and (3) corresponding reliance on arithmetical computation. Accordingly, the navigational chart is involved in the prototypical Western position-fix cycle of measurement, computation, and interpretation. Interpreted as a model of space, the chart suggests the travel concept of a series of chart locations (with the birds-eye perspective noted by Thorndyke and Hayes-Roth [75]).

Following Simon in *The Sciences of the Artificial*, Hutchins asserts that "[t]he fix cycle is accomplished by the propagation of representational state across a series of representational media" [37] in order to make the solution transparent. A representational state is a configuration of elements in a medium, which can be interpreted as a representation of something; representational states are carried between media by coordinating the states of the media. Broadly interpreted, computation can be described in the same way. In this model, an analog-to-digital conversion can facilitate easy information transportation and recording. In general, then, a cognitive system includes both internal and external processes, and communication; this model is thus neither mentalistic nor behavioral. Symbols are developed as internalized physical tools. (In this view, mental maps are internalized versions of physical ones.) According to Hutchins, tools don't amplify human cognitive abilities, but *they transform a task to a domain where a solution is evident*. Simple tools implement computation as simple manipulations of physical objects, and conceptual judgments as perceptual inferences. Social organization can be regarded as a computational architecture, "the simultaneous coordination of many media in a functional system that transcends the boundaries of individual actors" [37], in addition to coordination within an individual. Cognition is thus both individual and sociocultural. Learning can be seen as an adaptive reorganization of a complex system, even within an individual. Ultimately, Hutchins claims, the real power of human cognition is the ability to flexibly construct functional systems that accomplish our goals by bringing bits of structure into coordination.

Hutchins work has several implications for wayfinding research. First, it suggests the need for tools that transform a wayfinding task to a domain where a solution is evident (as studied in virtual wayfinding research [21], [70].) Hutchins research also recalls the diversity of Western navigational tools which could be adapted for electronic worlds. Because real space and time don't exist in these worlds, the representational assumptions underlying traditional tools must be carefully examined. (Some physical principles for virtual world design have been proposed [7].) Moreover, Hutchins' research suggests the power of multiple, specialized tools used in parallel or in series. (Researchers in information visualization have considered this possibility [48].) Finally, Hutchins' research shows the power of team navigation; computer-supported collaborative work (CSCW) may provide new techniques for online navigation.

3.5 Hypermedia Environments

Issues of information exploration arise in hypermedia as well as in spatial worlds. 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.* conducted research with text and hypertext documents ([49], [50]). 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 [25]. During learning, people successively instantiate schemata with landmarks, routes, and surveys [77]. 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" [49]. 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 review of research on spatial metaphors and disorientation in hypertext browsing is provided by Kim and Hirtle [42]. Much of this material is reviewed in the present paper. In addition, Kim and Hirtle identify three classes of disorientation problem:

- *navigational*
- *"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 - general wayfinding and task specific.

User navigation in complex database structures is characterized by Canter *et al.* In this work, the authors apply psychological concepts of navigation and path algebras [12]. This research considers account 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.* [79], and Ingram and Benford [40]). The authors' research suggests the value of characterizing users' navigational strategies in hypermedia such as the WWW; this characterization could be used to design Web sites to support the strategies associated with planned user tasks. The relative influence of user differences, task, and hypermedia topology is worth investigating. In this vein, Parunak identified the navigational strategies associated with particular hypermedia topologies [60]; this research will be discussed below. Finally, the navigational strategies of Canter *et al.* recall the optimal foraging models of Pirolli and Card [64]; further research might investigate this potential connection.

Focusing on differences between browsing and searching, Campagnoni and Ehrlich conducted a study of users' strategies and skills for IR in a hypertext help system [11]. 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 with an average range of system experience participated in the experiment. Subjects acted as system administrators to answer six user questions. Data were collected by software monitoring and videotape recording. After the session, each subject 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 noted that route knowledge is amenable to verbal description [84]. Second, semantic models could be proposed to users, rather than physical ones; this distinction is discussed below.

In a relatively large study, Catledge and Pitkow conducted an experiment to characterize browsing strategies in the WWW [16]. 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.*, serendipitous browsers are wanderers ; general-purpose browsers are scanners, browsers, and explorers; and searchers are searchers [12].) 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 studied users' revisitation patterns in Web navigation [75]. The researchers sought to gather empirical data for the design of history mechanisms in Web browsers. For six weeks, the browsing data of twenty-three users was gathered. Results showed that 58% of document accesses were revisits, and that users continually expand their document access set. Also, users often return to pages recently visited; they access relatively few pages often; they tend to browse in small clusters of related pages (a working set), and they repeat 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 [16] and the Tauscher and Greenberg [75] 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. Research is ongoing in this area (e.g., [5], [55]).

3.6 Discussion

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 internal and external structures. (This issue will be discussed below.)
- 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" and data collection mechanisms.
- Environment, tools, and individual differences affect user strategies and performance, to varying degrees.

4. INFORMATION STRUCTURING

The structure of an electronic world affects user navigation at both the articulatory and semantic levels. For locomotion, information structure determines the routes that a user can follow. For wayfinding, 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.

4.1 Structure Types

Adapting work of Rennison and Strausfeld in information visualization (discussed later in this review), we can identify five major types of information structure [67]:

- *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, a structure can often be implemented as a different structure. The categories of Rennison and Strausfeld provide a useful conceptual vocabulary for designing and discussing navigational research. Lynch's five city elements influenced and correspond to this vocabulary:

Urban Element (from Lynch)	Structure Type (from Rennison and Strausfeld)
path	graph
edge	linear
district	set
node	radial
landmark	relational

As noted by Benedikt, effective information structuring requires a selection of quantifiable attribute(s) from the data domain [7]. He proposed four types of attribute: geographical, chronological, alphabetical, and domain-specific. As noted by Wurman, domain-specific attributes can be subdivided in two ways: discretely as categories, or continuously as a continuum [85]. Rennison and Strausfeld suggest a natural correspondence between domain attributes and types of information structure [67]:

Attribute (from Benedikt, Wurman)	Structure (from Rennison and Strausfeld)
category	set
hierarchy	graph
geography	relational
continuum	linear
alphabetical	linear

This correspondence can be useful in the design of electronic worlds, especially for establishing a strong correspondence between semantic and physical structures [50].

Parunak discusses the theoretical implications of information structures for the problem of wayfinding in large hypermedia systems [60]. The author first identifies common navigational strategies in the physical world. He then defines a set of hypermedia topologies, each with appropriate navigational strategies. Parunak proposes that "appropriately restricting the connectivity of a hypermedia database can improve the ability of users to navigate" [60]. Accordingly, common hypermedia navigation tools may be seen as inducing a simpler topology on the hypermedia database. Two types of tool are common - "beaten path" mechanisms and typed links. "Beaten path" mechanisms include back-up stacks (e.g., history lists), path macros, and bookmarks. Typed links have specific advantages: "a hypermedia database structured as a set of distinguishable hierarchies will offer navigational and other cognitive benefits that an equally complex system of undifferentiated links does not . . ." [60]. Maps are also useful for large systems with constrained topologies.

Parunak's research [60]is a sort of dual to that of Canter *et al.* [12]. Canter *et al.* conducted a behavioral experiment in which user strategy is unconstrained by hypermedia topology; Parunak conducted a theoretical analysis in which user strategy is constrained by hypermedia topology. It appears that the navigational strategies employed by the behavioral subjects are composed of sub-strategies, which may conform to the more theoretical analysis. Future research might investigate Parunak's ideas behaviorally, as well as considering their adaptation to spatial worlds. Research might also consider wayfinding support for the navigational strategies implied by information structures. Parunak's comments corroborate Passini' emphasis on the route-finding role of maps [61], while recalling the unresolved question of hypermedia survey knowledge noted by McKnight *et al.* [49].

Regardless of the environmental structure or user strategy, navigation can be difficult in complex, large-scale information structures. Furnas studied this problem and developed some design recommendations [29]. 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 small view and short paths (small diameter) in the view graph. During design, inefficiently view-traversable structures may be improved by adding a traversable infrastructure (c.f., [15], [24], [83]). Furnas defines a link's *to-set* as the targets to which the link leads efficiently; an *inferred to-set* is a link's apparent to-set. A node's outgoing link information is defined to be well-matched if the node's to-set matches its inferred to-set with respect to all possible targets. A target's information residue or scent (per information foraging theory [64]) 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. Domain semantics should be used to subdivide an information structure efficiently.

Furnas offers a few examples of 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.

4.2 Hierarchical structures

Hierarchical structures feature prominently in research on navigation, information structuring, and information visualization. According to Furnas and Zacks, useful hierarchical structures fall on a continuum between trees and general graphs [31]. 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. Many hypermedia systems, including the WWW, appear to have a general graphical structure. Intermediate structures, DAGs naturally support semantic abstraction; like general graphs, though, DAGs suffer from a lack of constraints. (Another intermediate structure is the pretree [55].) 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 , who argues strongly for viewing and designing cities as "semilattices" (DAGs), rather than as trees [2]. 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" [2]. 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. Moreover, preliminary research by Modjeska and Marsh in the Web domain support Alexander's claim of user bias toward tree-like mental models [52]. In both

Alexander's and Furnas' models, multi-purpose elements are essential for successful design. In both domains, good structure may be more complex than designers can easily visualize; future research into design tools may be needed to solve this problem.

Alexander later extended the ideas about hierarchical urban structure to a general philosophy of design [3]. 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.

4.3 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 Norman *et al.* [58]) has been discussed in connection with mental maps and wayfinding. The gap between semantic and physical structures (the "articulatory distance" of Normal *et al.*) has been mentioned several times without detailed discussion. Several perspectives on this issue will now be reviewed.

McKnight *et al.* discuss the navigational roles of semantic intentions and physical representations [50]:

> 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 physical 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 [50].

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 and Furnas made similar recommendations for wayfinding [61] and effective view navigation [29], respectively.

In a related discussion, Kaplan and Moulthrop contrast physical and semantic "spaces" [41]. 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" [41]. In thought experiments, they propose egocentric weighted hypermedia links, as well as dynamically reconfigured hypermedia structure maps. Paraphrasing, the authors seek to display dynamic, local semantic information in hypermedia, in order to facilitate navigation.

The comments of McKnight *et al.* above [[50] recommend designing a stable environmental structure with a close correspondence between semantic and physical structures. A large, well-structured Web sites is an example of such an environment. The argument of Kaplan and Moulthrop, by contrast, suggests designing dynamic environmental cues and tools to insert semantic information into physical structures [41]. A large text collection with context-dependent, dynamic hypertext links is an example of such an environment [34]. These two sets of theoretical recommendations are complementary; research is needed to determine the relative importance and roles of the two approaches.

In related research, Shum discusses the properties and benefits of Euclidean (physical) and virtual (semantic) spaces [72]. 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. McKnight *et al.* proposed that physical structures reflect semantic ones [49]; Kaplan and Moulthrop proposed that dynamic, semantic information be represented locally in physical structures [41]; and Shum argues for the nesting of physical and semantic structures [72]. Again, further research is need to clarify the relative roles of physical and semantic structures in electronic worlds.

4.4 Experimental Results

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

Valdez *et al.* investigated the role of landmarks in user orientation during hypermedia browsing. The authors note the advantages of fisheye lens models [30] 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 [45] and Passini [61]; Valdez *et al.* confirm this finding for hypermedia. Semantically-enhanced landmarks would follow the recommendations of McKnight *et al.* [49] and Kaplan and Moulthrop [41]. 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 [54]. 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 obtained 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.

Preliminary research by Modjeska and Marsh has questioned the behavioral significance of structural landmarks computed by this algorithm [54].

Following the study by Valdez *et al.*, exploratory research was undertaken by Modjeska and Marsh on the WWW [52]. The study investigated the relationship between (1) site size and structure and (2) user navigation and perception. The authors conducted a behavioral experiment with structured Web browsing. A variety of data collection methods were used, including navigation logging, map sketching, interviews, and automated site analysis. Experimental results showed that strongly hierarchical sites tend to be more usable. In this study, site size has relatively little effect on user navigation and mental models. Nodes high in a hierarchy tend to be the most memorable ones. Surprisingly, strongly hierarchical sites are perceived as smaller than weakly hierarchical sites of corresponding actual size, other factors being equal. Results also showed that structural landmarks identified by the GIT algorithm [54] tend to be poor predictors of behavioral landmarks. The authors' results confirm the user preference for hierarchical mental models noted by Alexander [2] and Furnas [31]. Also, the authors' experimental methodology contributed to the relatively small body of research literature on Web navigational behavior (e.g., [16], [75]). Future research is needed in this area to consider a greater variety of Web site sizes and structural types.

4.5 *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 is generally credited with proposing the general concept [9]. 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. [9]

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.

A sort of implementation of the "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 [8]. The WWW is defined to include 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 recent estimates, the Web now has approximately 40,000,000 documents, 2,000,000 servers, and an unknown number of users [57]. Current users 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. 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 [7]. 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. [7]

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 [7]:

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) [7]. 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" [7]. Like Alexander [2] and Furnas [29], 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 on account 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 ([41], [49], [72]). Benedikt adopts a novel approach, proposing a 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 ([41], [49], [72]), but it is probably more computationally demanding as well.

4.6 Discussion

This section has reviewed some research on information structuring from a variety of disciplines - architecture, hypermedia, computer science, and psychology, among others. 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:
    - a taxonomy 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.

5. INFORMATION VISUALIZATION

As a consequence of the information explosion, the quantity of information to be assimilated for leisure and especially work has increased dramatically. The information to be assimilated falls broadly into two categories, identified by Benedikt as navigational and destinational [7]. The bulk of this review 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 data as well. (Besides, as mentioned above, multi-purpose elements can be designed to serve as both navigational and destinational data.) This section reviews key research in the field of information visualization: fisheye views, semantic zooming, dynamic queries, XEROX PARC's Information Visualizer, 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.

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

- hypermedia or spatial world
- type(s) of information structure
- technique for integrating physical and semantic structures
- value for wayfinding - mental mapping, route planning, or plan execution
- destinational or navigational data
- appropriate for electronic environments or maps
- representation of values or structure, which can be either raw or derived [79]
- presentation attributes, such as static or dynamic organization, simple or compound display, two or three dimensions, and textual or graphical presentation

5.1 Fisheye Views

A central problem in information systems is the lack of display space to show all available information. Furnas proposed generalized fisheye views as a solution to this problem [30]. 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, they often become lost and lack interpretative contextual information. Fisheye views solve this problem by trading off *a priori* 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 have proposed an extension of fisheye views for 2D graphs and other structures. They enhance the fisheye technique with layout considerations, including object positions, and detail size and level. These layout attributes are 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 have developed a prototype system, which can maintain real-time response for graphs containing a hundred vertices and a hundred edges. The authors note potential extensions to their work for multiple viewing foci and slave viewers for node content.

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 distorts the distance judgments needed for developing survey knowledge and cost-effective route plans in spatial worlds. The technique displays data which is both navigational and destinational. For navigation, the technique may be more appropriate for maps than for environments, since dynamic fisheye distortions may interfere with the user's mental mapping.

5.2 Semantic Zooming

Another graphical technique to balance detail and context 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 have recently proposed a general analytic framework for multi-scale interfaces [32]. Their "space-scale diagrams" represent a spatial world and its possible magnifications, which allows the analysis and visualization of related UI issues.

The Pad project developed a prototype interface based on semantic zooming [62]. 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 [23]. 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 treats links as detail items into which to zoom.

Bederson and Hollan designed a successor to Pad, named Pad++ [6]. 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 propose zooming as part of an "interface physics", which uses physical models for the visible behavior of objects (as suggested by Benedikt [7]). 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 shows 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 ([19], [61]. [73]). The technique displays data that is both navigational and destinational. For navigation, the technique is suited to both environments and maps, particularly the latter [6]. The technique can be used in two- or three-dimensional worlds.

5.3 *Dynamic Queries*

Accessing information in databases is a major activity of knowledge workers. Unfortunately, traditional database query languages trade off ease of use for power. To solve this problem, Ahlberg, Williamson, and Shneiderman proposed the technique of dynamic queries [1]. This technique is a convenient visualization of local database queries, with a simple, intuitive, interactive query refinement method. The technique uses direct manipulation ([39], [70]) to construct queries. One or more selectors (UI widgets) control the value range of one or more attributes. Viewing a graphical database representation, users manipulate the selectors to explore data subsets rapidly and easily. In a navigational environment, dynamic queries may offer a useful way to reveal attributive information, which can facilitate wayfinding.

Ahlberg and Shneiderman have recently enhanced a version of dynamic queries [1]. Their design goals include "rapid filtering to reduce result sets, progressive refinement of search parameters, continuous reformulation of goals, and visual scanning to identify results" [1]. The authors developed several tools:

- dynamic query filters – GUI widgets to adjust parameters
- starfield displays -- 2D scatterplots that support semantic zooming
- tight coupling - direct manipulation; coordinating query components with display invariants; and progressive query refinement

Software prototypes include a geographical display of housing availability and a starfield display of film database information. The authors report the tools to be natural and effective for simple queries.

Related research in the database lab at the University of Toronto produced a set of tools for dynamic querying and visualization [51]. Combining a sort of dynamic query with hypermedia structure visualization, one tool displays 2D network graphs using a variety of algorithms. Small histograms below a graph reveal additional attributes of the data. User 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 supports a sort of semantic zooming through sub-graph nesting to arbitrary depth. The system features a visual query language, which uses the same elements as the visualization: "queries are annotated (hy)graph patterns whose matches on the database can be displayed in different ways" [51]. A slave Web document browser allows the user to move between structural and document views. (A slave browser is also used in [23] and [83].)

A UI tool derived from dynamic queries and semantic zooming, the MagicLens filter was developed at XEROX PARC [74]. Intuitively, a MagicLens is a semantic magnifying glass for examining on-screen information. More formally, a MagicLens combines a display region with an information operator to change the view of objects in the region. Filters can be composed, or clicked through to manipulate objects. Filters have several advantages over fixed applications: an intuitive physical metaphor, modified views in the original context (a sort of fisheye view), clutter limited to a small region, easy compound query construction, and an extensible and portable paradigm. Useful filters include the following types:

- *local views* - displaying details in context, reducing clutter, multiple simultaneous views, or queries with built-in geometric attributes;
- *coordinated alternate views;*
- *visual macros;*
- *safe document modification previews;*
- *enhanced editing operations.*

In informal testing, the movable filters were found to be usable and appealing.

A team of Japanese researchers have combined dynamic queries with multiple visualizers and semantic zooming [48]. Their work shows the power of coordinating multiple visualization styles and tools for effective navigation. The team created a prototype of a tourist information system for the city of Nara. The WING system (Whole Interactive Nara Guide) tightly couples four views. Direct manipulations in each visualizer (un)constrain the data displayed in other visualizers:

- The graphical 3D *map view* supports zooming and viewpoint movement as a simple VR system, which displays major city features.
- The *content view* shows pictures and as a travel guidebook. Items nearest the center of the map view are displayed at the top of the content list.
- The *category view* organizes all data items hierarchically, as a semantically zoomable 2D text grid.
- The *index view* lists data names. The list is a permuted index, in which entries appear multiple times, sorted by substrings. For searching, users magnify and shrink the gap between adjacent terms, as a lexicographical zoom. Informal testing showed good user performance with this view.

Dynamic query techniques are suitable for hypermedia or spatial worlds, with many information structures. They can display existing values and structure (original dynamic queries and WING), gather values and derive structure (Web visualizer), or handle values and structure according to technique type (MagicLens). Functionally, physical structures handle the input and output of semantic actions. Formally, the visualizers of different techniques relate physical and semantic structure in different ways. In the house locator and Web structure visualizer, for example, semantic information labels physical structure. In the WING system, physical and semantic structures exist in adjacent containers. In a MagicLens filter, finally, semantic zooming reveals either nested physical structure or parallel semantic structure. Dynamic queries are likely to facilitate mental mapping, by automating the coordination of locational and attributive information [25]. The technique is also suited for route planning, since it could be used to procure required information on a timely basis. Dynamic queries can display data that is destinational (original dynamic queries and MagicLens) or both navigational and destinational (Web visualizer and WING) . The technique is most appropriate for electronic maps, since dynamic modifications of an environment might interfere with wayfinding, unless presentation changes preserved essential structure and relationships. As described above, dynamic queries incorporate a range of presentation attributes.

5.4 XEROX PARC: The Information Visualizer

This section will review theoretical principles and six visualization designs of a research project named the Information Visualizer at 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 XEROX PARC have developed an information workspace known as the Information Visualizer (IV) ([14], [65]). The IV provides a framework for several visualizations that incorporate computer graphics , systems architecture, and cognitive psychology. A key design principle is "that many parts of a task can be off-loaded from the user's cognitive system onto their perceptual system, thus enabling more rapid processing. [65]. 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 workspace can generally be described as "a special environment in which the cost structure of the needed materials is tuned to the requirements of the work process using them" [14]. The goal of the IV project, then, is to develop a workspace for information.

Following these goals, the IV has three architectural components:

- The capacity of the user's immediate storage (memory) is increased by a technique called *3D/Rooms*. Rooms features non-immersive VR as a container for information visualizers. 3D graphics technology maximizes the effective use of screen space, enabling the display of large information structures.
- The rate of the user-system interaction and information transfer is increased through an architecture termed the *Cognitive Coprocessor.* This architecture uses UI animation, so that human perception can track structural relationships without conscious thought. The Cognitive Coprocessor features multiple asynchronous agents and smooth, interactive animation. Display quality is reduced in minor ways during motion, to achieve quick visual transitions. The system is tuned for three human time constants:

    - perceptual processing (0.1 seconds)
    - immediate response (1 second)
    - unit task (10 seconds)
- *Information visualizations* are structure-oriented browsers of information sets. They increase the level of information abstraction. Three types of information are visualized - unstructured, linear, and hierarchical. Several browsers use “focus+context” techniques, which are abstracted fisheye views.

Note that the Cognitive Coprocessor architecture has several implications for navigation in electronic worlds:

- A Cognitive Coprocessor architecture would help users to maintain context during navigation.
- Increasing the bandwidth of user-system communication may improve navigation, particularly with a well-structured environment and tools.
- Shifting some wayfinding cognition to the perceptual system could free resources for information and management tasks [42].

PARC researchers have made several observations about information work, which inform the IV’s visualizations:

- Users often refine queries iteratively, respecifying questions and improving techniques during a task session (c.f. dynamic queries [1]).
- Users often access multiple, heterogeneous sources. They use meta-information about these sources to facilitate their work. Important categories of meta-information are often automatically extractable. These categories include content, provenance, form, functionality, and usage statistics.
- Users often interleave access in a parallel fashion, when using sources with slow or variable response. Accordingly, a good user model includes feedback on time costs and the status of ongoing operations.
- Users usually mix access with other parts of the work process. For example, they mix searching or browsing with visualization or analysis. Multiple visualizations can facilitate the integration of such an information workspace (c.f. the WING prototype [48]).

A focus+context visualizer of linear information, the Perspective Wall uses a physical metaphor to fold wide 2D layouts into 3D visualizations [47]. A center panel shows detail, while two wing panels in perspective view show context. Any 2D task-specific features are retained. The technique is appropriate for work information spaces with spanning properties such as time. The information may be layered, e.g., for simultaneous projects. This Perspective Wall solves two important problems

- displaying a large volume of information
- fitting an extreme aspect ratio on screen

This technique shows a three-fold improvement in information quantity and visual size over full-width 2D presentations. Sample domains include file systems and corporate documents.

Another focus+context technique, the Cone Tree assists with managing and accessing large hierarchical information spaces [47]. Cone trees are hierarchies laid out in 3D. 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.

The hyperbolic browser is a radical fisheye visualizer for large hierarchical structures [44]. Mathematically, the technique lays out a hierarchy on a uniform hyperbolic plane, and then maps this plane onto a 2D circular display region. 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.

Butterfly is a specialized, compound IV application for accessing DIALOG's Science Citation databases on the Internet [46]. 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. Simple 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.

For access to Web information, IV researchers propose two visualizations [15]. The WebBook, first, is a 3D interactive book of HTML pages. It allows rapid interaction with information at a higher level of aggregation than individual pages. In this visualization, 3D graphics and interactive animation show relationships between pages. The books support flipping through pages; bookmarks; scrollbars; and an exploded, fisheye view. The partly serial nature of books is thereby enhanced with other electronic forms. WebBooks can be assembled by either the user or the system. The book metaphor was chosen for its useful properties:

- As in a library, Web information is often heterogeneous, with clusters existing in loosely structured environments; such clusters often have a sequential organization.
- Books use display space efficiently.
- Books and their affordances (properties for use) are familiar to users.

The WebBook is an instance of a navigable superstructure for a complex underlying structure, as recommended by Furnas [29].

The Web Forager embeds WebBooks in a hierarchical 3D workspace. It allows users to interact intensively and rapidly with large numbers of Web pages. In terms of interaction rates, the virtual workspace is arranged as a three-level hierarchy:

1. a focus place for direct user-content interaction - a large book or page
2. an immediate memory space for objects in use, but not the immediate focus - the air and a desk
3. a tertiary place for storing many pages and books - a bookcase

The immediate storage place has several tiers by z-distance in 3D. The workspace can simultaneously display many objects. It thereby supports tradeoffs on a continuum between fast access to few objects, and slow access to many objects. Following information foraging theory [64], the Web Forager supports rapid access to high-value information patches.

All of the IV tools use non-immersive VR, except for the hyperbolic browser. Each tool is specialized for a particular information structure. While the Perspective Wall, 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.* [49]). The display information is simultaneously navigational and destinational, although only the Butterfly and Web Forager have well-differentiated places. For this reason, the visualizers are better suited for use as tools than as environments. Nevertheless, 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 and Web Forager for cost-benefit analyses, and asynchronous agents for timely information-gathering. Several tools could also facilitate plan execution, particularly the Perspective wall and WebBook, 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.

5.5 Landscapes for Information Retrieval

Another visualization approach involves information landscapes for 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 [45] and Passini [61]. The prototype Bead system had several design goals, including a legible information space [45], reliance on sensori-motor skills [14], 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. The system uses multi-dimensional scaling to determine document proximities. 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 [45]. 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. Ongoing research includes improved navigation.

In later research, Chalmers, Ingram, & Pfranger explored the addition of imageability (legibility) features to the Bead system[17]. 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; 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 (virtual personae), but it visually represents the activities of all current users. (Related earlier work is reported in [40].)

LyberWorld, another spatial visualization of a document corpus, supports IR through navigation in complex, abstract information spaces [34]. 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 are placed in correspondence. In Bead, the landscape shows 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 demonstrates a successful adaptation of physical principles for electronic worlds. It also shows 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 captures an aspect of real-world experience, as well as showing a potential social aspect of electronic navigation.

5.6 Spaces for Public Information

Relatively static information spaces are often complex and not effectively navigable. This comment applies particularly to public electronic worlds, e.g., the WWW. 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 [29], a Web visualizer with graphical query support [51], and a 3D information workspace [15]. 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 developed a prototype that combines textual VR with a Web browser ([23] [24]). He proposes that spatial UI metaphors facilitate user navigation in hypertext, by introducing location, distance, and direction. Dieberger's Juggler system is a so-called "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. [29]). The imagined environment is organized as a semi-regular grid. MOO locations can have three types of exit:

- (local) *directional exit*
- (local) *non-directional exit*
- (non-local) *special exit* - "magic" feature that functions as a hypertext link (e.g., teleporter).

Combinations of such exits can support three types of navigational topology:

- *Euclidean* - requires linear navigation
- *Euclidean with exceptions* - has parallel spaces for navigation
  - *Euclidean* - mediated by textual semantics, for browsing familiar information
  - *semantic* - mediated by physical hypertext, for finding new environmental routes
- *non-spatial topology* - has semantic structure

Juggler uses a Euclidean topology with exceptions (c.f. [41]). The new Euclidean structures offer a reduced, more navigable version of Web structure (c.f. [60]). 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 ([81], [83]). 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 proposal for Information Islands has two components -- a hierarchically-structured public world and a mobile, private environment [81]. 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

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 developed two parallel design sketches of spatial worlds for the WWW [83]. “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” [83]. 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. In the likely absence of a unified cyberspace, following such recommendations may be important. Note the similarity of StackSpace to the Web Forager [15] in supporting an emergent, spatial Web superstructure, with a strong use of the depth display dimension. Web Forager features superior specialization of display regions, but StackSpace offers greater flexibility.

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 ([19], [73]) and route planning [61], 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.

5.7 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 approaches 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, and navigational issues can be challenging. Research in this area will be reviewed in this section.

The seminal work hypermedia structure visualization the SemNet project [27]. 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 attempted:

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. [15], [83]).

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. The first two of these techniques are powerful and may be combined. 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" [27]. The quality of the 3D imagery is essential (as in XEROX PARC's Information Visualizer [14]). 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 have adapted Lynch's legibility research [45] to enhance of existing visualizations [40]. 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 Mosaic browsing, and manual node repositioning. Ingram and Benford report generally positive experiences. Future goals include better 3D path navigation and improved general navigational tools. (Related later work is reported in [17].)

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 visualizations are compelling, and informal usability testing has been positive.

The File System Navigator, a popular hypermedia structure visualization, is distributed with Silicon Graphics (SGI) computers. This application uses a landscape metaphor to display the 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, and 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 ([4], [5]). 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.

Although more general, Harmony's tightly coupled navigational and visualization tools (mixing graphics and text) recall the tools of the WING prototype [48].

Researchers in the Graphics, Visualization, and Usability Laboratory at the Georgia Institute of Technology (GIT) have developed the Navigational View Builder (NVB), a navigational visualization of complex hypermedia networks such as the WWW ([55], [56]). 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 both the translation of network to a hierarchy, and the visualization of the resulting hierarchy.

The hypermedia viewers discussed in this section all handle graphical structures, whether tree (FSN), directed graph (SemNet), or general graph (LEADS, Narcissus, Hyper-G, and Navigational View Builder). 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.

5.8 Other Systems

This section will review three research projects in information visualization that don't fit conveniently in one of the categories above. The first project explores navigation in a dynamic, semantically-structured space. The second, related project investigates the relationship between human navigational metaphors and information system structures. The third project considers nested, Cartesian worlds for financial visualization.

At MIT's Media Lab, Rennison has developed the Galaxy of News system [66]. A prototype USENET news system, Galaxy draws on fisheye views, cone trees, and Pad to support user navigation in an abstract, multidimensional information space. The author's design goals include 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 smoothly 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 principles:

- 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.

Given the disorienting effects of Web navigation in unfamiliar domains, Rennison may enhance Galaxy for Web support in the future. Note that the Galaxy system is customized for a DAG information structure. The system displays a world with some spatial properties - motion, exclusive location, and distance - and some semantic properties - no coordinate axes or origin, no fixed positions, and relative, changing, multi-valued relationships. Physical structures continually parallel semantic ones, as far as possible in an electronic instantiation. Accordingly, the world has some navigational properties, including context, target-seeking, and movement . But wayfinding in a normal sense is not possible, because the world is dynamic, multiply structured, and non-Euclidean. Displayed information is both semi-navigational and semi-destinational. The system is thus too specialized to support an effective, general information world; but it could be an effective, evocative information tool.

In related research, Rennison and Strausfeld developed the Millennium Project, “a conceptual and computational approach for enabling understanding of a large, multidimensional set of information” [67]. The system allows the user to explore a semantically structured space by navigation and manipulation. The project helps the user to understand relationships between information 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 [43] 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. The system’s computational process has three steps:

1. analyzing the information base to build structural representations
2. presenting the information in a 3D virtual space with appropriate contexts and structures
3. interpreting user movements and actions, for continuous querying and dynamic reconstruction of the virtual space.

This research offers a foundation for the semantics of objects and user actions in virtual worlds. As suggested by Benedikt [7], virtual worlds can benefit from the assignment of semantics to behavior modeled on the physical world. Benedikt discussed primarily external experience (e.g., motion); Rennison and Strausfeld discuss primarily internal experience (e.g., orientation). The two approaches seem complementary for wayfinding design in virtual worlds.

Feiner and Beshers considered information that is structured with high, fixed, intrinsic dimensionality, e.g., financial data [28]. Such data can be displayed using techniques of information design or, using 3D graphics, as nested information worlds. These researchers have developed a software prototype for interactive, financial visualization in *n* dimensions. The system’s new interaction metaphor is called "worlds within worlds": "nested heterogeneous coordinate systems that allow the user to view and manipulate functions" [28]. In creating an inner coordinate frame, the origin's position in the outer frame indicates the values of the outer variables, which are held constant within the inner frame. Using the a VR Data-Glove, frames can be translated (to reset outer variables within the inner frame) or rotated (for better viewing). The order of variable assignment is significant in determining the final visualization. The system's UI features three tools:

- the *dipstick* - for inspecting embedded values
- the *waterline* - for inspecting embedded lines
- the *magnifying box* - a sort of 3D magnifying glass

Feiner and Beshers' spatial worlds handle compound linear (Cartesian) structures of high dimensionality. In a way, these worlds constitute a sort of immersive zooming. Physical and semantic structures are placed in correspondence, and displayed data is both navigational and destinational. The dynamic nature of the visualization hinders long-term wayfinding. Otherwise, the system has the three wayfinding elements recommended by Passini [61]: a strong organizational scheme (including position, distance, and direction); good spatial enclosure (each coordinate frame); and good spatial correspondence (precise and clear nesting). The system also incorporates versions of three Lynchian features [45]: edges (frame borders), landmarks (frame corners and key data points), and districts (frames). So short-term wayfinding is well supported. The visualization has the problem of maintaining user orientation in deeply nested contexts, as well as locational ambiguity at the edge of a coordinate frame. Despite certain difficulties, the technique of nested worlds seems generalizable to contexts beyond financial visualization. Benedikt, for example, suggested a similar scheme for an architectural database [7].

5.9 Discussion

A diversity of visualization techniques have 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 have 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' recommendations for navigable superstructures in complex structures [29].

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., the Information Visualizer [14]), and the opposite approach that seeks to generate strong cognition through basic perceptual cues (e.g., the Juggler model [23]). 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 psychological techniques are both required for effective navigation in electronic worlds.

6. CONCLUSION

This paper has reviewed research on navigation in electronic worlds. Three major research areas have been discussed: exploration, structuring, 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 more important than 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: psychologically-oriented UI architectures; appropriate physical and 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.

In general, there are many types of navigational task, electronic environment, user strategy, navigational tools, and individual user. Consequently, there are many opportunities for research and design in hypermedia and virtual worlds. A short summary includes the following opportunities:

- a taxonomy of tasks, environments, and strategies
- a taxonomy of current navigational tools and designs for new ones (e.g., dynamic electronic maps)
- the adaptation and development of structuring and visualization techniques for electronic worlds

- the role of non-visual sensory modalities, e.g., audio or kinesthetic cues
- the nature and role of temporal cues (i.e., hyper-time and virtual time)
- implications for social navigation, CSCW, and virtual communities
- interactions between types of real and virtual environment for learning, browsing, and other tasks

Research in these areas will be valuable for understanding the human activity of navigation. Such research will be essential for designing and developing usable, large-scale electronic worlds.

7. ACKNOWLEDGMENTS

I would like to acknowledge my thesis advisors, Mark Chignell and Marilyn Mantei, for their suggestions and guidance during the writing of this paper.

8. BIBLIOGRAPHY

Navigation in Electronic Worlds:
Research Review for Depth Oral Exam

David Modjeska
Department of Computer Science
1 May 1997

1. Christopher Ahlberg and Ben Shneiderman, "Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays." In *Human Factors in Computing Systems: Proceedings of the CHI '94 Conference*. New York: ACM, 1994.

2. Christopher Alexander, "A City is Not a Tree." In John Thackara, ed., *Design after Modernism: Beyond the Object*. London: Thames and Hudson, 1988.

3. Christopher Alexander, *The Timeless Way of Building*. New York: Oxford University Press, 1979.

4. Keith Andrews, Frank Kappe and Hermann Maurer, "Serving Information to the Web with Hyper-G." In *Computer Networks and ISDN Systems. Proceedings of the Third International WWW Conference*. 1995.

5. Keith Andrews, "Visualizing Cyberspace: Information Visualization in the Harmony Internet Browser." In *Proceedings of IEEE Information Visualization 1995*. Washington: IEEE, 1995.

6. Jan Bedersen and J. Hollan, "Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics." In *Proceedings of UIST '94*. New York: ACM, 1994.

7. Michael Benedikt, "Cyberspace: Some Proposals." In Michael Benedikt, ed., *Cyberspace: First Steps*. Cambridge, Massachusetts: The MIT Press, 1991.

8. Tim Berners-Lee, Robert Cailliau, Ari Luotonen, Henrik Frystyk Nielsen, and Arthur Secret, "The World-Wide Web." In *Communications of the ACM*. New York: ACM, 1994.

9. Bertin, *Graphics and Graphic Information Processing*. Berlin: deGruyter Press, 1977.

10. Vannevar Bush, "As We May Think." In *The Atlantic Monthly*. 1945.

11. Campagnoni and Kate Ehrlich, "Information Retrieval Using a Hypertext-Based Help System." In *ACM Transactions on Information Systems*. New York: ACM, 1989.

12. David Canter, Rod Rivers, and Graham Storrs, "Characterizing User Navigation through Complex Data Structures." In *Behaviour and Information Technology*. 1995.

13. David Canter, "Cognitive Cartography." In *The Psychology of Place*. London: Architectural Press, 1977.

14. Stuart K. Card, George G. Robertson and Jock D. Mackinlay, "The Information Visualizer, an Information Workspace." In *Human Factors in Computing Systems: Proceedings of the CHI '91 Conference*. New York: ACM, 1991.

15. Stuart K. Card, George G. Robertson William York, "The WebBook and the Web Forager: An Information Workspace for the World-Wide Web." In *Human Factors in Computing Systems: Proceedings of the CHI '96 Conference*. New York: ACM, 1996.

16. Lara D. Catledge and James E. Pitkow, "Characterizing Browsing Strategies in the World Wide Web." In *Computer Networks and ISDN Systems 27*. Elsevier Science, 1995.

17. Matthew Chalmers, Robert Ingram and Christoph Pfranger, "Adding Imageability Features to Information Displays." In *Proceedings of UIST '96*. 1996.

18. Matthew Chalmers, "Using a landscape metaphor to represent a corpus of documents." In Andrew U. Frank and Irene Campari, eds., *Spatial Information Theory: A Theoretical Basis for GIS: European Conference, COSIT '93*. Springer-Verlag, 1993.

19. William G. Chase, "Spatial Representations of Taxi Drivers." In R. Rogers and J. A. Sloboda, eds., *Acquisition of Symbolic Skills*. New York: Plenum, 1983.

20. William G. Chase, "Visual Information Processing." In K.R. Boff and J.E. Lincoln, eds., *Handbook of Human Perception and Performance*. New York: Wiley and Sons, 1986.

21. Rudolph P. Darken and John L. Silbert, "A Toolset for Navigation in Virtual Environments." In *Proceedings of ACM User Interface Software and Technology (UIST) '93*. New York: ACM, 1993.

22. Rudolph P. Darken and John L. Silbert, "Wayfinding Strategies and Behaviors in Large Virtual Worlds." In Michael J. Tauber, ed., *Human Factors in Computing Systems: Proceedings of the CHI '96 Conference*. New York: ACM, 1996.

23. Andreas Dieberger, "Browsing the WWW by interacting with a textual virtual environment -- A framework for experimenting with navigational metaphors." In *Proceedings of ACM Hypertext '96*. New York: ACM, 1996.

24. Andreas Dieberger, "Providing Spatial Navigation for the World Wide Web." In Andrew U. Frank and Werner Kuhn, eds., *Spatial Information Theory: A Theoretical Basis for GIS: International Conference COSIT '95*. Berlin: Springer, 1995.

25. Roger M. Downs and David Stea, "Cognitive Maps and Spatial Behavior: Process and Products." In Rober M. Downs and David Stea, eds., *Image and Environment*. Chicago: Aldine, 1973.

26. Thomas D. Erickson, "Working with Interface Metaphors." In Brenda Laurel, ed., *The Art of Human-Computer Interface Design*. Reading, Massachusetts: Addison-Wesley, 1990.

27. Kim M. Fairchild, Steven E. Poltrock and George W. Furnas, "SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases." In Raymonde Guindon, ed., *Cognitive Science and its Applications for Human-Computer Interactions*. Hillsdale, NJ: Erlbaum, 1988.

28. Steve Feiner and C. Beshers, "Worlds within Worlds: Metaphors for Exploring n-Dimensional Virtual Worlds." In *Proceedings of UIST '90*. New York: ACM, 1990.

29. George W. Furnas, "Effective View Navigation." In *Human Factors in Computing Systems: Proceedings of the CHI '97 Conference.* New York: ACM, 1997.

30. George W. Furnas, "Generalized Fish-eye Views." In *Human Factors in Computing Systems: Proceedings of the CHI '86 Conference*. New York: ACM, 1986.

31. George W. Furnas and Jeff Zacks, "Multitrees: Enriching and Reusing Hierarchical Structure." In *Human Factors in Computing Systems: Proceedings of the CHI '94 Conference*. New York: ACM, 1994.

32. George W. Furnas and Benjamin B. Bederson, "Space-Scale Diagrams: Understanding Multiscale Interfaces." In *Human Factors in Computing Systems: Proceedings of the CHI '95 Conference.* New York: ACM, 1995.

33. Gene Golovchinsky, personal communication, 1997.

34. Gene Golovchinsky, "Queries? Links? Is there a difference?." In *Human Factors in Computing Systems: Proceedings of the CHI '97 Conference.* New York: ACM, 1997.

35. Matthias Hemmje, Clemens Kunkel and Alexander Willett, "LyberWorld -- A Visualization User Interface Supporting Fulltext Retrieval." In W. Bruce Croft and C.J. van Rijsbergen, eds., *SIGIR '94: Proceedings of the 17th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*. London: Springer-Verlag, 1994.

36. Hendley, N. S. Drew, A. M. Wood and R. Beale, "Narcissus: Visualising Information." In *IEEE Visualization '95*. IEEE, 1995.

37. Edwin Hutchins, *Cognition in the Wild*. Cambridge, Massachusetts: MIT Press, 1995.

38. Edwin Hutchins, "The Technology of Team Investigation." In Jolene Galegher, Robert E. Kraut, and Carmen Egido, eds., *Intellectual Teamwork*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1990.

39. Edwin L. Hutchins, James D. Hollan and Donald A. Norman, "Direct Manipulation Interfaces." In Donald A. Norman and Stephen W. Draper, eds., *User Centered System Design: New Perspectives on Human-Computer Interaction*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1986.

40. Rob Ingram and Steve Benford, "Legibility Enhancement for Information Visualization." In *Proceedings of IEEE Visualization '95*. 1995.

41. Nancy Kaplan and Stuart Moulthrop, "Where No Mind Has Gone Before: Ontological Design for Virtual Spaces." In *European Conference on Hypermedia Technology 1994 Proceedings*. New York: ACM, 1994.

42. Hanhwe Kim and Stephen C. Hirtle, "Spatial Metaphors and Disorientation in Hypertext Browsing." In *Behaviour and Information Technology*. Taylor & Francis, 1995.

43. George Lakoff and Mark Johnson, *Metaphors We Live By*. Chicago: The University of Chicago Press, 1980.

44. John Lamping, Ramana Rao, and Peter Pirolli, "A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies." In *Human Factors in Computing Systems: Proceedings of the CHI '95 Conference*. New York: ACM, 1995.

45. Kevin Lynch, *Image of the City*, The. Cambridge, Massachusetts: The MIT Press, 1960.

46. Jock D. Mackinlay, Ramana Rao andStuart K. Card, "An Organic User Interface for Searching Citation Links." In *Human Factors in Computing Systems: Proceedings of the CHI '95 Conference*. New York: ACM, 1995.

47. Jock D. Mackinlay, George G. Robertson and Stuart K. Card, "The Perspective Wall: Detail and Context Smoothly Integrated." In *Human Factors in Computing Systems: Proceedings of the CHI '91 Conference*. New York: ACM, 1991.

48. Toshiyki Masui, Mitsuru Minakuchi, George R. Borden IV, and Kouichi Kashiwagi, "Multiple-View Approach for Smooth Information Retrieval." In *Proceedings of UIST '95*. New York: ACM, 1995.

49. Cliff McKnight, Andrew Dillon, and John Richardson, *Hypertext in Context*. Cambridge, England: Cambridge University Press, 1991.

50. Cliff McKnight, Andrew Dillon, and John Richardson, "Space--the FInal Chapter or Why Physical Representations Are Not Semantic Intentions." In Cliff McKnight, Andrew Dillon, and John Richardson, eds., *Hypertext: A Psychological Perspective*. New York: Ellis Horwood, 1993.

51. Alberto Mendelzon, "Visualizing the World Wide Web." White Paper. Toronto: University of Toronto, 1996.

52. David Modjeska and Anna Marsh, "Structure and Memorability of Web Sites." Working Paper in the Department of Industrial Engineering. Toronto: University of Toronto, 1997.

53. Joy Mountford, "Tools and Techniques for Creative Design." In Brenda Laurel, *The Art of Human-Computer Interface Design*, 17-30. Reading, Massachusetts: Addison-Wesley, 1990.

54. Sougata Mukherjea and James D. Foley, "Showing the Context of Nodes in the World-Wide Web." In *Human Factors in Computing Systems: Companion to the CHI '95 Conference*. New York: ACM, 1995.

55. Sougata Mukherjea, James D. Foley, and Scott Hudson, "Visualizing Complex Hypermedia Networks through Multiple Hierarchical Views." In *Human Factors in Computing Systems: Proceedings of the CHI '95 Conference*. New York: ACM, 1995.

56. Sougata Mukherjea and James D. Foley, "Visualizing the World-Wide Web with the Navigational View Builder." In *Computer Networks and ISDN Systems 27*. Elsevier Science, 1995.

57. Jakob Nielsen, *Multimedia and Hypertext: The Internet and Beyond*. Boston: AP Professional, 1995.

58. Donald A. Norman, "Cognitive Engineering." In Donald A. Norman and Stephen W. Draper, eds., *User Centered System Design: New Perspectives on Human-Computer Interaction*. Hillsdale, New Jersey: Erlbaum, 1986.

59. Donald A. Norman, “Some Observations on Mental Models.” In Dedre Gentner and Albert L. Stevens, eds., *Mental Models*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1983.

60. Van Dyke Parunak, “Hypermedia Topologies and User Navigation.” In *Proceedings of the Hypertext ’89 Conference*. New York: ACM, 1989.

61. Romedi Passini, *Wayfinding in Architecture*. New York: Van Nostrand Reinhold, 1984.

62. Ken Perlin and David Fox, “Pad: An Alternative Approach to the Computer Interface.” In *Proceedings of the SIGGRAPH ’93 Conference*. New York: ACM, 1993.

63. Peter Pirolli, “Computational Models of Information Scent-Following in a Very Large Browsable Text Collection.” In *Human Factors in Computing Systems: Proceedings of the CHI ’97 Conference.* ACM: New York, 1997. pp. 3-10.

64. Peter Pirolli and Stuart Card, “Information Foraging in Information Access Environments.” In *Human Factors in Computing Systems: Proceedings of the CHI ’95 Conference*. New York: ACM, 1995.

65. Ramana Rao, Jan O. Pedersen, Marti A. Hearst, Jock D. Mackinlay, Stuart K. Card, Larry Masinter, Per-Kristian Halvorsen, and George G. Robertson, “Rich Interaction in the Digital Library.” In *Communications of the ACM*. New York: ACM, 1995. pp. 29-39.

66. Earl Rennison, “Galaxy of News: An Approach to Visualizing and Understanding Expansive News Landscapes.” In *UIST 94, ACM Symposium on User Interface Software and Technology*. New York: ACM Press, 1994.

67. Earl Rennison and Lisa Strausfeld, The Millenium Project.” In Andrew U. Frank and Werner Kuhn, eds., *Spatial Information Theory: A Theoretical Basis for GIS: International Conference COSIT ’95*. Berlin: Springer, 1995.

68. George G. Robertson, Jock D. Mackinlay, and Stuart K. Card, “Cone Trees: Animated 3D Visualizations of Hierarchical Information.” In *Human Factors in Computing Systems: Proceedings of the CHI ’91 Conference*. New York: ACM, 1991.

69. Manojit Sarkar and Marc H. Brown, “Graphical Fisheye Views of Graphs.” In *Human Factors in Computing Systems: Proceedings of the CHI ’92 Conference*. New York: ACM, 1992.

70. Glenna A. Satalich, “Navigation and Wayfinding in Virtual Reality: Finding the Proper Tools and Cues to Enhance Navigational Awareness.” M.S. Thesis. Seattle: University of Washington, 1997.

71. Ben Shneiderman, “Direct Manipulation.” In *Designing the User Interface*. Reading, Massachusetts: Addison-Wesley, 1992.

72. Simon B. Shum, “Real and virtual spaces: mapping from spatial cognition to hypertext.” In *Hypermedia*. London: Taylor & Graham, 1990.

73. Stevens and P. Coupe, “Distortions in judged spatial relations.” In *Cognitive Psychology*. 1973.

74. Maureen C. Stone, Ken Fishkin and Eric A. Bier, "The Movable Filter as a User Interface Tool." In *Human Factors in Computing Systems: Proceedings of the CHI '94 Conference*. New York: ACM, 1994.

75. Linda Tauscher and Saul Greenberg, "Revisitation Patterns in World Wide Web Navigation." In *Human Factors in Computing Systems: Proceedings of the CHI '97 Conference.* New York: ACM, 1997.

76. Perry W. Thorndyke and Sarah E. Goldin, "Spatial Orientation: Theory, Research, and Application." In *Spatial Learning and Reasoning Skill*. Plenum Press: New York, 1983.

77. Perry W. Thorndyke and Barbara Hayes-Roth, "Differences in Spatial Knowledge Acquired from Maps and Navigation." In *Cognitive Psychology*. New York: Academic Press, 1982.

78. Edward C. Tolman, "Cognitive Maps in Rats and Men." In Herbert S. Langfeld, ed., *The Psychological Review*. Lancaster, Penn.: American Psychological Assoc., 1948.

79. Lisa Tweedie, "Characterizing Interactive Externalizations." In *Human Factors in Computing Systems: Proceedings of the CHI '97 Conference.* New York: ACM, 1997. Pp. 375-382.

80. Felix Valdez, Mark Chignell, and Bernice Glenn, "Browsing Models for Hypermedia Databases." In *Proceedings of the Human Factors Society, 32nd Annual Meeting*. 1988.

81. John A. Waterworth, "Information Islands: Private Views of Public Places." In *Proceedings of MHVR '94: East-West Conference on Multimedia, Hypermedia and Virtual Reality*. Moscow: 1994.

82. John A. Waterworth and Mark H. Chignell, "A Model for Information Exploration." In *Hypermedia*. 1991.

83. John A. Waterworth, "Personal Spaces: 3D Spatial Worlds for Information Exploration, Organisation and Communication." In *Proceedings of BCS Conference on 3D and Multimedia on the Internet, WWW, and Networks*. Bradford, UK: 1996.

84. Christopher D. Wickens, "Spatial Perception and Cognition and the Display of Spatial Information." In *Engineering Psychology and Human Performance*. New York: Harper Collins, 1992.

85. Richard Saul Wurman, *Information Anxiety*. New York: Bantam, Doubleday, Dell, 1989.