

How Spatial Is Hyperspace? Interacting with Hypertext Documents: Cognitive Processes and Concepts

PATRICIA M. BOECHLER

ABSTRACT

The World Wide Web provides us with a widely accessible technology, fast access to massive amounts of information and services, and the opportunity for personal interaction with numerous individuals simultaneously. Underlying and influencing all of these activities is our basic conceptualization of this new environment; an environment we can view as having a cognitive component (hyperspace) and a social component (cyberspace). This review argues that cognitive psychologists have a key role to play in the identification and analysis of how the processes of the mind interact with the Web. The body of literature on cognitive processes provides us with knowledge about spatial perceptions, strategies for navigation in space, memory functions and limitations, and the formation of mental representations of environments. Researchers of human cognition can offer established methodologies and conceptual frameworks toward investigation of the cognitions involved in the use of electronic environments like the Web.

INTRODUCTION

FOR CENTURIES, TEXT-BASED LITERATURE has been the primary mode for disseminating and receiving information. The cognitive processes involved in reading and writing books have been studied extensively for decades as they have played a crucial role in the development of humankind. In the past two decades, hypertext systems (computer-based information systems of interconnected documents) have become increasingly pervasive in work, home, and educational environments. With the advent of the World Wide Web, these information systems are heavily influencing how we develop as a society. However, unlike traditional text information sources, the World Wide Web is interactive. This has created not

only a new cognitive environment ("hyperspace") where information is perceived, stored, manipulated, and retrieved in new ways, but also a new social environment ("cyberspace"), where one individual's cognitions about this new environment interact with those of other people who share it. Unlike physical space, the World Wide Web is not constrained by distance. Interactions between individuals can be numerous and instantaneous. The possibilities for enhancing individuals' lives through easily accessed personal interactions and the availability of online services are abundant.

The study of hypertext systems is a multidisciplinary venture. Psychologists, computer scientists, and educators all bring different approaches and goals to the effort. Cognitive psychologists wish to know how humans concep-

tualize hyperspace and the nature of the mental representations they form of hyperspace. Social psychologists want to understand how this new way of communicating impacts an individual's sense of self, their social sphere, the group dynamics of this new environment, and the effect on society at large in both a regional and global sense. Clinical psychologists are investigating the ways that these systems can be used to provide psychological services. Computer scientists and software developers want to know how to design an optimal interface for users. Educators are interested in how this technology can be applied to enhance learning. This can be advantageous as a broad knowledge base and diverse perspectives are brought to bear on the topic. However, it also results in the lack of a unified theoretical foundation for hypertext. Relative to many other areas of psychological study, research on the cognitive aspects of using such systems is in its infancy; models and paradigms are still being developed.

One point to note is that the above-mentioned endeavors are not as distinct as one might initially suppose. Basic cognitive processes (e.g., memory, perception, etc.) are the underpinnings for more complex processes (e.g., social-emotional responses to our environment). In the real world, we've identified many relationships between basic cognitions and social responses.^{1,2} There are numerous instances of our perceptions of physical space producing psychosocial consequences (e.g., personal space being violated, overcrowding, feeling lost). It is reasonable to suppose that how we perceive hyperspace may affect our social interactions within it. Understanding our cognitive reactions to hyperspace may well aid us in predicting some of our social and emotional responses to this new environment. To fully take advantage of the opportunities offered by cyberspace, we need to comprehend clearly how individuals take in information from hyperspace.

To this end, the traditional paradigms and theories of cognitive psychology have much to offer, providing us with established methodologies and conceptual frameworks for exploring different facets of cognition relevant to hypertext navigation, such as memory, task

performance, and spatial processing. Understanding the cognitive aspects of hyperspace will help us to develop optimal technologies that will guide the way we use the social environment of cyberspace. Understanding the cognitive processes that shape many different individuals' perceptions of hyperspace will facilitate the creation of technologies that are accessible to a wide range of people with a wide range of characteristics (e.g., age differences) and skills (e.g., level of literacy). For the uninitiated, traversing the World Wide Web in search of solutions to everyday problems can be extremely daunting and can often end in failure to locate useful information or access needed services. Providing navigation tools that are based on the cognitive processes of users will do much toward making the World Wide Web an accessible resource for everyone.

Psychological processes are often described via metaphors.^{3,4} Hyperspace is a dynamic, interactive space requiring new psychological metaphors, new ways of thinking and interpreting information. However, before we can devise new metaphors that portray this environment accurately, it is prudent to closely examine previous metaphors to identify beneficial aspects that can act as a foundation for new metaphors.

One overarching metaphor seems to dominate the hypertext literature: the spatial metaphor. The spatial metaphor arose out of a need for a common language to discuss hypertext issues and a framework within which to develop usable interfaces. Its basic premise is that locating information in hyperspace has similar psychological features to navigating in physical space. Therefore, promoting mental representations of spatial layouts of information is a good approach to improving the hypertext user's ability to access information.^{5,6} Navigational tools modeled after real-world navigational aids such as maps appear to enhance performance on some measures⁷⁻¹¹ but not on others.¹²⁻¹⁶ Correlational data on the relationship between spatial ability and performance has shown increases in speed of performance with higher spatial ability,^{17,18} but marginal effects on the accuracy of performance.¹⁹ Therefore, the empirical evidence to date has not resolved the issue of how or to

what degree spatial processing is involved in the efficiency and effectiveness of hypertext use.

The spatial metaphor seems easily understood by novice users. However, if it is not an accurate metaphor for the mental processes that humans engage in as they interact with hypertext, or if we do not understand the aspects of the comparison that make the metaphor beneficial, then as more and more complex, powerful systems are developed this metaphor will eventually be less useful. This would be so for both the goal of achieving optimal access to information and the goal of presenting information in a form that can enhance social interaction or learning.

The following review draws from many diverse areas of cognitive psychological study to bring forth questions whose answers could greatly impact the benefits we are able to garner from cyberspace. It is an in-depth look at the possible ways we can view the new environment of the World Wide Web. This may or may not be as a space, therefore, this review also critically examines the spatial metaphor for hypertext interaction.

AN INTRODUCTION TO HYPERTEXT

What is hypertext?

Generally, a hypertext system is an electronic system of interconnected units of information or text nodes. Many different types of hypertext systems have been designed, each consisting of its own specific structures that support the connections (links) between the user interface (the windows or pages a user views on the display terminal) and the underlying information system (database of text nodes). The user interface is sometimes referred to as the "front end" of a hypertext system and the underlying database as the "back end" of the system. In this paper, the term "window" will be used in reference to the screen image a user may be viewing and "node" in reference to the informational node that the window is connected to in the underlying database. The computer environment of hypertext documents has come to be referred to as "hyperspace." Hypertext doc-

uments that contain images, audio, and video clips are referred to as "hypermedia" documents. In the hypermedia literature the terms "hypertext" and "hypermedia" are often used interchangeably.²⁰

Each hypertext window contains icons that can be pictorial symbols or highlighted words (hotwords). A window can contain any number of these icons. The icon represents the link to the database for that particular piece of information. When clicked on with a pointing device (usually a mouse), the icon activates the link to the referenced node in the database and a new window is instantaneously opened containing the information in that node. Nodes that hold related information are linked. Some nodes will have a multitude of links, both to and from other nodes based on their relatedness, whereas other nodes may only have incoming links. It is important to stress that links are not merely connections between nodes but that they should specify the semantic relationships between nodes. The author of the hypertext document is the engineer of the links based on his or her semantic network; that is, his or her conceptions of the interconnectedness of ideas within the chosen material.^{21,22}

Links can be organized in two ways: in a hierarchical structure or in a network structure. In a hierarchical structure, nodes are only linked to the superordinate information above them or the subordinate information below them.^{21,23} In a network structure, any node can be linked to any other. Multiple links can exist between superordinate and subordinate information.²³ Regardless of the specific structure, because of its general node/link relational structure, a hypertext system is nonlinear²⁴; that is, unlike the pages in a book, users can access different windows in different orders. There is no predetermined sequential organization of the underlying nodes. Typically, a hypertext system will also have a backtrack function to allow the user to access previously viewed windows.²⁵

A brief history of hypertext systems

The concept of hypertext was first outlined by Vannevar Bush in 1945 with his Memex informational system.²⁴ This was a mechanical

system rather than an electronic system. It consisted of several microfilm projection units on which the viewer could display different frames of microfilm. The viewer could construct links, which Bush called trails, between two items by viewing the two frames simultaneously, then inserting code that denoted the names of the frames and their current positions within each document. The important feature of the Memex was that any item within the system could be made to immediately and automatically access any other item. The key idea was that two items could be linked together for fast and flexible access. The system itself was never implemented, but Bush had the foresight to envision a time when his idea could be applied to computer systems, which at the time were in the early stages of development, inaccessible and costly.

Since that time, many versions of hypertext systems have been developed and continue to be developed. Each system has its own specific features, but Bush's initial idea is still the most essential feature; that the system consists of a viewer connected to underlying informational nodes that can be accessed with speed and flexibility by the user.

In 1963, Engelbart devised the oN Line System (NLS), with features we would easily recognize as attributes of current hypertext systems.²⁴ The NLS highlighted three particular aspects of hypertext: "a database of non-linear text, view filters which select information from this database, and an interface which structured the display of this information for the terminal." (p. 22).²⁴ The accessibility of high-resolution displays prompted the focus on graphical displays of nodes (the single page display we are so familiar with now) and the one-to-one correspondence of node to window. Engelbart also devised the "mouse." Engelbart's system focused on the front end of a hypertext system, the user interface.

Concurrently, Ted Nelson was developing a UNIX-implemented system that focused on the back end of a hypertext system, the database server. He called this system "Xanadu." The goal of the Xanadu project was to develop a system capable of storing and retrieving a massive corpus of literary works, a system that was easily expandable through the inclusion of new

links and windows rather than through some central, and cumbersome, change in the system structure. The term "hypertext" was implemented by Nelson in 1965.²⁴

Since the initial implementation of these systems, a multitude of hypertext systems have been developed, each with its own purpose and specialized features. Systems have been developed to address very definite informational purposes, and general hypertext systems have been devised for the purpose of studying hypertext itself as a technology.

The variations of function and features within different hypertext systems are so numerous and complex as to be well beyond the scope of this paper. What is important to note is that as these systems evolved, grew in size and intricacy, and became accessible to the general public (culminating in the early 1990s with the advent of the World Wide Web), the necessity for understanding how humans interact with hypertext, and perceive and manipulate information from the system, became increasingly important.

To summarize, the defining features of hypertext are: a large collection of informational nodes connected via machine-supported links, which support massive cross-referencing and easy access. The term "hypertext" will be used as a general term for systems that entail these features regardless of their individual differences.

Goals in hypertext use

Different hypertext users have different goals, which produce different navigational strategies. Users may be searching out a very specific piece of information that they have already accessed, looking for a specific piece of information that they do not know how to find, or they may be exploring the system with no specific informational goal. Conklin²⁴ outlines the three methods of conducting an informational search in a hypertext system: (1) the user can visit a set of related windows or pages through the traversal of links; (2) the user can navigate through the document using a graphical browser, which provides a graphical representation of the network structure; or (3) the user can activate a queried search where the

system searches all documents to find any instance of a keyword or key phrase that is entered. The query method is necessary for documents or networks of documents, such as the World Wide Web, that cannot be exhaustively examined because they contain too many nodes.

Advantages and disadvantages of hypertext

The main advantage of hypertext over traditional text is that it is a high-speed method of accessing vast quantities of information. As mentioned earlier, there is no predetermined sequential organization of the underlying nodes. This provides an increased flexibility in the ways that users can experience the information in the document compared to traditional text-based presentations. The user is free to make link connections in any order based on their own focus and interpretations of the material, toward their own end goal. However, this same feature that provides such individual freedom for the user also comes with costs. For example, at any point in time, the user only sees the current window and the links represented by icons or highlighted text, leading from that window to other related windows. This means the user must remember the windows already viewed and hypothesize about the overall size, structure, and complexity of the entire hypertext network, as it can never be viewed in its entirety.

Because of this, there are two related categories of problems associated with using hypertext systems: cognitive overhead and disorientation.

Cognitive overhead. Unlike traditional text, hypertext does not have conventional cues that lead the reader through a document (e.g., chapters, sections).²⁶ With traditional text documents, format often indicates the type and purpose of a document and gives the reader hints about where important information will be found (e.g., a legal contract is easily distinguished from a grocery list). In hypertext, no such conventions have been established. In addition, there are multiple routes or paths to access a specific piece of information in the network. This increased flexibility results in users

being responsible for decisions about which information to access (which links to select) and in what order. To this end, the cognitive resources of the user must be allocated toward several goals concurrently. The term “cognitive overhead” has been used to refer to the amount of cognitive resources necessary to successfully complete an informational task in hypertext.²⁴ With hypertext, users must evaluate information currently viewed, must recall information previously viewed as well as its location relative to other viewings, and must create a plan to navigate toward desirable information (not only deciding which windows to access but also which tools to use to execute the plan). Kim and Hirtle⁶ label these cognitive tasks as (1) navigational tasks: planning and executing routes through the network; (2) informational tasks: reading and understanding the contents presented in the nodes and their relationships, for summary and analysis; and (3) task management: coordinating informational and navigational tasks. The cognitive demands that these tasks require are compounded by increasingly complex and voluminous hypertext structures.^{22–24,27,28} Degraded performance occurs when a hypertext user’s cognitive resources are overwhelmed by these demands.

Disorientation. One of the symptoms of cognitive overhead is disorientation.^{24,28,29} It is well documented that hypertext users will often report “feeling lost” when performing poorly on an informational search.^{22,26,29,30} Based on subjective reports, Edwards and Hardman³⁰ characterized this phenomenon via three categories of the users’ experience: (1) the user does not know where to go next; (2) the user knows where to go but not how to get there; (3) the user does not know where they are in relation to the overall structure of the document. Based on empirical measures, Foss³¹ proposed three types of disorientation problems and their symptoms: (1) Navigational disorientation problems, caused by a lack of knowledge of organization and extent of the document as well as unfamiliarity with navigational tools. The symptoms of navigational disorientation are looping, inefficient pathways and query failures. (2) The embedded digression problem caused by high cognitive de-

mands that lead to difficulties in planning, managing and executing digressions. The symptoms for embedded digression are a disorganized screen layout, having many windows open concurrently and engaging in excessive backtracking. (3) The art museum problem caused by the high cognitive demand of viewing many windows superficially without attention to detail or relatedness. The symptoms of the art museum problem are short reading times and restricted search paths.

The problems of disorientation and cognitive overhead were identified by early hypertext researchers,^{22,23,29,30} but continue to be unresolved.^{13,28} It has been suggested that unless disorientation and cognitive overhead can be effectively reduced the ultimate usefulness of hypertext will be severely limited.^{24,28,29}

The spatial metaphor

The identification and labeling of the two problems of disorientation and cognitive overhead are a direct result of how we have come to conceive of the environment that hypertext creates. Disorientation is a cognitive phenomenon that occurs in an environment where it is necessary to maintain a sense of one's own location relative to the location of other objects and to maintain the relative location of these objects to each other. Cognitive overhead in hypertext is the amount of cognitive resources that are required to successfully interact with the type and number of structures within that environment (the type of structures dictated by our conception of the environment).

The spatial metaphor arose out of a need for a common language to converse about the hypertext environment and, subsequent to that, out of the need for a framework for designing tools and hypertext features that create optimal navigating environments. Canter et al.⁵ have been attributed with the emergence of the metaphor¹³ which has been advocated by a number of researchers.^{6,8,9,15,32} However, other authors have expressed concern over the rapid and seemingly overwhelming acceptance of the metaphor.^{13,33-36}

Having presented an overview of hypertext systems, and a brief description of how users interact with them, it should be clear how

firmly entrenched the spatial metaphor is in the hypertext literature. The use of terms related to space is solidly embedded in the topic of hypertext. Terms such as "hyperspace," "environment," and "navigation" are difficult to avoid as, at this point in time, they provide for the user the most direct access to meaning at an introductory level. To this end, the spatial metaphor has been useful and may continue to be useful. However, beyond the introductory level, is the spatial metaphor impeding the development of a more cognitively accurate model, one that might ultimately provide more benefit to the user?

Before addressing this larger issue, it is prudent to attempt to make explicit the ways in which the terms "spatial" and "metaphor" are used (or not used) in relation to hypertext.

WHAT IS THE ROLE OF A METAPHOR?

The research on metaphor and its related term, analogy, is vast and complicated. The following section is a brief review of perceptions of metaphor in both the psychological and human-computer interaction literature.

In the psychological literature, Lakoff³⁷ describes metaphor as a mode of thought that can be understood as a tightly structured mapping from the source domain to the target domain. The source domain is a conceptual schema that we have constructed based on our experiences in a particular environment. In the case of the spatial metaphor, the source domain is physical space and the target domain is hyperspace.

Holyoak and Thagard³ consider metaphor, like analogy, a method toward understanding a new domain in terms of a source domain, but metaphor involves connections that move beyond surface category structure. "Many metaphors are based on deeper relational and system mappings" (p. 217).³

Holyoak stresses the importance of the evaluation of these mappings based on similarity, structure, and purpose between the source and target domains.

According to Pylyshyn⁴ the role of a metaphor for scientific explanation is not to fully explain all aspects of a complex environment but to provide a framework on which to reference

new, vague, and disconnected ideas about that environment, phenomenon, or concept. This framework is a means toward the development of a more thorough understanding of something new. Pylyshyn describes the acquisition of new knowledge using Piaget's terms of accommodation and assimilation. When we first encounter new information, we attend to aspects that are most similar to the schemas we've already constructed and categorize the new information as belonging to the most similar schema (assimilation). If the new information is so dissimilar that it cannot be mapped onto an existing schema, we then adjust and expand the structure of the schema to include this new knowledge (accommodation). However, Pylyshyn maintains that it is important to provide some degree of explanation via the metaphor (a powerful metaphor) rather than merely description, which may leave the user with a groundless impression of having been provided with an explanation (an impotent metaphor). The value of a metaphor decreases if the user, upon closer inspection, realizes that the metaphor has few shared properties with the new situation the person is attempting to understand. Surface similarities may prove of limited value as functional and conceptual dissimilarities are uncovered.

In hypertext, as the user becomes more proficient the comparison between hyperspace and physical space may break down. The spatial metaphor may be useful as a pedagogical tool to introduce the novice user to the system and to aid in the initial organization of his or her thoughts about this new environment, but may also impede the user from developing a more in-depth and explicit conceptualization of the hypertext environment. As Pylyshyn notes,⁴

One must distinguish between the general programmatic enterprise of trying to illuminate a new phenomenon and the much more demanding goal of establishing the validity of an explanatory theoretical principle. (p. 548)

Therefore, one goal in metaphor use within psychology is to define and explain the relations between the source and the target domain in the metaphor, and the cognitive processes involved in executing the task associated with that metaphor. It seems the spatial metaphor

has been prematurely accepted as a theoretically plausible metaphor for the cognitive processes involved in hypertext interaction.

Unlike psychological theory, in the human-computer interaction (HCI) literature the end goal is to provide a metaphor for the user that helps them understand the purpose and function of the facilities of a specific application, rather than to provide the user with an explanation of the global environment of hypertext. Metaphors can be applied to an application with a narrow purpose (e.g., the typewriter metaphor for word-processing) or a broader purpose (e.g., the general desktop metaphor used in many types of applications). In the HCI literature, some authors are careful to explicitly present the spatial metaphor as a design tool for a specific application. An example of this is Hammond and Allinson's³² travel holiday metaphor that includes choices for go-it-alone travel or guided tour travel for a specific hypertext system. Other authors present the spatial metaphor as if they accept the assumption that it's a psychologically plausible account of how people think about hyperspace in general.^{6,8}

Interface designers expect and accept the presence of metaphorical mismatches. Carroll et al.³⁸ emphasize the usefulness of mismatches as do Hammond and Allinson.³² Hammond and Allinson contend that a metaphor need not be accurate at all levels of description to be useful:

Furthermore, a metaphor which is self-evidently inappropriate at some levels of description will be more successful than one which is ambiguous: no attempts will be made to invoke mappings at the inappropriate level. (p. 83)

An example of a metaphor mismatch is exhibited in the metaphor of word processor as typewriter. On a typewriter, the space bar moves the typing position, whereas on a computer keyboard the space bar inserts a space. Novice users are often confused when the spacebar does not produce the same result on a computer that it does on a typewriter.³²

Although mismatches are expected, these researchers also advocate a more systematic assessment of metaphor to interface design. Carroll et al.³⁸ argued that three aspects of analysis

must be undertaken to properly assess the value of an interface metaphor: operational, structural, and pragmatic analyses. Operational analysis entails the measurement of behavioral effects (e.g., improvements in performance and learning). Structural analysis involves making explicit the mappings of referent domain to target domains, a description of primitives and relations between primitives in the referent and target domains. Structural analysis is necessary to determine the scope of the metaphor. Pragmatic analysis examines the practical context of the metaphor. This includes examining the user's goal associated with using the metaphor as well as the impact on the user of the flaws in the metaphor (e.g., mismatches). This is similar to Holyoak's³ focus on similarity, structure, and purpose in a psychological approach to metaphor validation. Carroll et al.³⁸ emphasized that all three types of analyses are necessary to fully understand the dynamics between metaphor and user and, hence, to assess the overall usefulness of the metaphor.

An adequate analysis of interface metaphors needs to rest on an empirical task analysis of what users actually do, not an abstract normative analysis of what they might do. (p. 73)

Psychological and HCI researchers espouse similar types of metaphor assessment but toward a different focus; respectively, understanding underlying information processing mechanisms versus usability. Both the psychological and HCI approaches seem plausible for the evaluation of a broad metaphor like the spatial metaphor. Nonetheless, in a survey of the literature, it seems neither approach has been used in a systematic way to closely examine the hypertext/real-space comparison. The mappings between hypertext features and real-world features have not been empirically specified nor have particular cognitive processes involved in the use of hypertext (e.g., degree of spatial processing) been identified in detail.

RELATING SPACE TO HYPERTEXT

While investigating the use of a spatial metaphor, it is also important to consider definitions of space itself. For example, what are

spatial properties and how do these properties map onto "hyperspace?"

Metric models of space

Real world, everyday, physical space is generally described in terms of geometric properties, specifically Euclidean or metric properties. Three-dimensional space is denoted by the geometric relationships between three points, X, Y, and Z. These relationships include such properties as direction and distance. Metric space adheres to three geometric axioms: minimality, symmetry, and triangle inequality.³⁹ These axioms will be described in detail presently.

Stanton and Baber³⁵ express concern that this everyday definition of space has been inappropriately applied to hyperspace (which they refer to as "electronic space"). They state:

... the basic concept of hypertext in computer science terms, is an n -dimensional space which can be traversed by moving through links. 'Space' as used in this context has a well defined meaning, as the collection of objects and activities contained within a specific domain, similar to its use in finite state architecture. However, much current research in hypertext appears to use the term 'space' in its everyday sense, that is, as a physical relationship between objects. (pp. 235–236)

In the case of hypertext, it appears that physical distance is being equated to some degree with psychological distance. What exactly is psychological distance? In hyperspace, one way of conceptualizing psychological distance would be an estimate of the degree of semantic relatedness or similarity between two windows.

Geometric models, like the Euclidean model of physical space, are often used in the psychological literature to describe psychological processes. Beginning in the early 1960s, many psychological researchers^{40–43} asserted that the similarity of objects could be modeled using geometric properties. More explicitly, the similarity of two objects could be represented as points in some coordinate space, and that a metric distance function could be used to compute the dissimilarity between two points. A metric distance function assigns a distance to

each pair of points based on adherence to the same three axioms that define physical space: minimality, symmetry, and triangle inequality. The minimality axiom implies that the similarity between an object and itself is the same for all objects. The assumption of the symmetry axiom is that the distance between two points is the same no matter the order; that is, the distance between A and B is the same as the distance between B and A. The triangle inequality axiom assumes that the shortest distance between two points is a straight line. That is, if A, B, and C are any three points in space, then the sum of any two distances (e.g., AB, BC, AC) is greater than or equal to the third distance. Applied to similarity theory, it implies that if A is similar to B, and B is similar to C, then A and C cannot be very dissimilar from each other.⁴⁴

Contrary to the above view, Tversky⁴⁴ and Tversky and Gati⁴⁵ argued that although similarity relations may be represented using geometric models including the axioms presented above, these representations should not be viewed as psychological theory. Tversky contended that the minimality axiom is not fulfilled by some types of similarity measures. For example, the probability of judging two identical stimuli as the "same" rather than "different" is not constant for all stimuli. In recognition experiments of simple objects, sometimes an object is identified as another object more frequently than it is identified as itself. According to Tversky, the symmetry axiom does not hold for similarity relations either. Similarity judgments can be regarded as extensions of similarity statements; that is, statements of the form "a is like b." In such a statement, a subject and a referent are evident. The statement is directional as it begins with the subject and ends with the referent. To reverse the statement is to lose its meaning; "b is like a" is not the same as "a is like b."

As for the triangle inequality axiom, Tversky stated that the implication mentioned above (if A is similar to B and B is similar to C, then A and C cannot be very dissimilar from each other) sets limits on the similarity between A and C, as the similarity between A and B, and B and C constrains the similarity between A and C. Tversky provided an example that chal-

lenges the psychological validity of the triangle inequality assumption. He makes comparisons between the following pairs of countries, Jamaica–Cuba, Cuba–Russia, and Jamaica–Russia. The dimension that defines similarity between Jamaica and Cuba (geographic similarity) is not the same as the dimension that defines the similarity between Cuba and Russia (political similarity, at the time). Therefore, assigning a degree of similarity between Jamaica and Russia is uninformative, as comparisons are not made on dimensions that are valid between the two countries.

How well do geometric properties map onto hyperspace? In hyperspace, distance could be thought of as the number of steps between windows (e.g., the number of mouse clicks), the number of layers to traverse in a hierarchical system, or the amount of time it takes to access one window from another.³³ However, it could be said that these variables are only epiphenomena of the underlying semantic relationships of the material as perceived by the author, the person who has created the structure of the hypertext document. The links between windows are constructed by the author based on his or her perceptions of the relatedness of meaning between the content of each window. In this sense, the links are somewhat arbitrary in that the degree of relatedness perceived by the author may be very different from the degree of relatedness perceived by the user. Also, depending on the order in which the user has accessed certain windows, these variables may not consistently represent, from user to user, the underlying semantic relationships between windows.

The question therefore becomes, in the hypertext context, how does meaning map onto space? One implication of the spatial metaphor is that the degree of semantic similarity between windows can be psychologically represented as a distance. Windows that are similar in meaning are "close together" in the "information space," windows that are not are "far" apart. This, in turn, implies that semantic similarity can be described using geometric properties, just like Euclidean geometric properties define physical space, hence, connections between windows can be conceptualized as "space." Although a "metric" distance is not

applied to each pair of hypertext windows, because of the comparison to real space prompted by the spatial metaphor, there seems to be an underlying assumption that some sort of distance relations exist.

However, as outlined above, Tversky and Gati^{44,45} provide some compelling reasons for why psychological estimates of similarity should not be represented geometrically (that is, as a metric distance) for all stimuli in all situations. The hyperspace environment seems to violate all the axioms necessary for a geometric similarity model. The minimality axiom is not always accurately adhered to with simple stimuli. Referring to semantic similarity between hypertext windows, this problem is intensified due to the complexity of each window and the sheer number of windows being accessed in large documents. The symmetry axiom is also violated. Once a new window is accessed, the information the user encodes there goes with him or her back to the previous window, where that window is likely to be viewed in a slightly different manner. Therefore, knowledge-wise, the initial "distance" from the first window to the second or subsequent windows is not the same as the "distance" from the second window back to the first. The "distance" has been altered merely through the accessing of the second window and its contents. As for the triangle inequality axiom, different windows may or may not share the dimensions that other pairs of windows share just as, in Tversky's example, Jamaica, Cuba, and Russia do not share the same dimensions.

This discussion of similarity theory has been applied toward the validity of using a geometric model for semantic similarity in hypertext just as a geometric model is used to describe real space (Euclidean properties). The depth of comparison in this example may be excessive, but it is a demonstration of how unspecified the parameters of the spatial metaphor remain. We are still left with an ill-defined and vague sense of how space is related to the hypertext environment and the elements within it.

Spatial primitives

A less formal way of defining space is through decomposition of its characteristics

into a set of basic primitives. Primitives can be thought of as both elements and relationships that are necessary to convey the fundamental attributes of a phenomenon. Golledge⁴⁶ identifies and describes four spatial primitives: (1) identity, (2) location, (3) magnitude, and (4) time. Golledge describes providing identity "as the process of equating an occurrence with a name or label. The purpose behind this differentiation is to allow occurrences to be recognized and their uniqueness evaluated" (p. 104). As Golledge notes, a problem with identity is that, at times, two labels may be given to the same entity (e.g., a street is known by different names to different people). In hypertext, if an identity is not universally acknowledged, not only may different users be unable to recognize an item as the same item but also individual users may confuse different items.

The second primitive, location, is denoted as "information about where an occurrence exists within the totality of the environment" (p.104). The way location is specified can be different depending on the environment being considered. The spatial terms and reference points that describe an entity's location will be different if that entity is a continent rather than a country, a city rather than a building or an object. Locational information includes information that describes relations between an entity and other entities within that specific environment. Understanding the idea of location in hypertext has several obstacles. How do we describe different categories of environments (e.g., what is global, what is local)? How do we specify location within that environment? An item in hyperspace has no absolute location or even a consistent relative location.

From the property of location, several derived concepts emerge: (1) distance, (2) direction, (3) sequence and order, and (4) connection and linkage. Golledge⁴⁶ defines distance as the "interval between the locations of occurrences" (p.106). To quantify distance, location must be clearly specified. Direction can only be specified within an established frame of reference. Sequence and order are dependent on definitions of distance and direction as well. Connection and linkage are related to principles of proximity and similarity which, again, are inextricably tied to the specification of location.

The relation between hyperspace and these derived concepts cannot be determined until the relation between hyperspace and location is clearly understood.

The third primitive, magnitude, refers to the size of an identity. In instances where we can define "how much" with a single number, the measurement of magnitude is straightforward (e.g., how many cows are in a field). In other instances, Golledge concedes that magnitude is difficult to measure. For example, do we define a city by the area covered, the size of the population, or by the boundaries of specific activities? With hypertext, do we measure magnitude by the number of nodes in a document, the degree of linkage within the document, or the degree of linkage to other documents within a network? It is not clear how magnitude may be defined in hyperspace.

The fourth primitive, time, denotes the permanency of an occurrence. Does an occurrence exist at a later point in time just as it exists in this moment? Does an occurrence retain its location over time? The relation between time and location is particularly tenuous in hypertext as we may not even be able to denote location itself. As well, large systems of hypertext documents, such as the Web, are dynamic. Using a particular site as a landmark for locating other sites will be ineffective as sites can be removed, linking can be changed, and the general appearance of the site may be altered. Beyond the issue of location, unlike real space, navigation time in hypertext is not correlated to "distance." Changing locations between windows that are unrelated can be achieved just as quickly as moving between windows that are closely related.

Whether we approach the hypertext/real space comparison from the axioms of Euclidean space or a list of spatial primitives, identifying discrepancies between the two environments will inform us about the value of the spatial metaphor toward enhanced information access.

Spatial cognition

If spatial cognition plays an important role in how people interact with hypertext, comparisons between behaviors in hyperspace and

behaviors in physical space should be possible on several levels: Do people hold spatial representations in memory of the hypertext environment? Do novice hypertext users go through the same process that people in novel physical environments do? Do children show similar patterns of development in hypertext navigation as real world navigation? To address these questions we must have a brief look at spatial cognition and how it develops in real-world environments.

The study of spatial cognition is characterized by diverse approaches in the choice of settings and paradigms. For example, subjects have been asked to find their way,⁴⁷ to locate an object in real space,⁴⁸ as well as to imagine how an object might look if it was rotated to a different orientation.⁴⁹ Small-scale environments such as a single room where all contents are visible have been studied⁵⁰ and, conversely, large environments, such as the layout of a campus or city, where all locations are not within view.⁵¹ Researchers have been interested in people's knowledge of particular places as well as their knowledge of abstract spatial concepts. Subsequently, there are a variety of definitions of spatial cognition, each prompted by and constrained by the environment under investigation.

In general terms, spatial cognition is the knowledge and mental representation of the structure of space and the location and relation of objects or entities within that space, along with the knowledge of when and how to use this stored information and the ability to think about the mental representations themselves.⁵² This definition includes two parts: the informational system itself and the individual's ability to access and make use of that information. The informational system has been referred to as a cognitive map, a term first used by Tolman.⁵³ In this paper, the term cognitive map will be used to refer to the informational system that is part of and supports spatial cognition.

Cognitive maps. Cognitive maps are long-term memory representations of spatial information. As we cannot view cognitive maps directly, we must infer their features from behavior during specific tasks. Consequently,

empirical evidence involves these representations as they are applied to the performance of a particular action. External forms produced by behavior (e.g., sketch maps, verbal protocols, quantitative and structural analyses), are often called spatial products, a term originally used by Liben.⁵² Spatial products give us some indication of the nature of underlying spatial representations. Behavioral evidence suggests what the rules and operations are that govern the use of spatial representations.⁵⁴

The term cognitive map often elicits the assumption that this representation is an image-based, map-like mental representation. The actual form of cognitive maps has been highly debated. Some researchers have suggested that the cognitive map is analogous to a cartographic map, a Euclidean model of the world.⁵⁵ This is not to say that cognitive maps are exact replications of features on a map. They have been described as "schematic, sketchy, incomplete, distorted, and otherwise simplified and idiosyncratic."⁵⁶

Other researchers have suggested a conceptual-propositional, rather than image-based, theory of spatial information encoding.⁵⁷ In this view, both verbal and visual information is stored as abstract conceptual propositions. Concepts are embedded in a propositional network and conceptual meanings are partially dependent on their relations to other concepts in the network. If concepts are close together in the propositional network, they will function as cues for the retrieval of one another. Alternatively, some researchers^{57,58} have argued that the term "map" should only be thought of as a hypothetical construct for the spatial system that we accept as existing and influencing our spatial decisions.

Aside from the debate about the form of cognitive maps, Kitchin⁵⁸ summarizes the types of information that cognitive maps contain. Cognitive maps are selective, subjective representations of the environment that preserve the locational and attributional features, including subjective emotional responses to those attributes, that are most informative for the individual perceiving that environment. The attributes that designate the structure of the map (landmarks or labels) are markers for action that should take place at that point in space and

time. According to Kitchin, this is an important aspect of the cognitive map, that it is constructed to combine the concepts of place and action. Cognitive maps are individual constructions for individual events: there is no singular cognitive map that retains all the spatial information and knowledge we've acquired. They are dynamic; that is, our knowledge (cognitive map) influences the experiences we have and our experiences alter our knowledge (cognitive map).

More recent research into the nature of cognitive maps has suggested that their hierarchical structure is important. Hierarchies are organized by semantic and conceptual relatedness, not by spatial relations. Several studies provide evidence that cognitive maps have a hierarchical structure; that is, entities are encoded at different levels or regions (e.g., particular objects in relation to sets of objects) based on their conceptual similarity or dissimilarity rather than their actual spatial location.

For example, Stevens and Coupe⁵⁹ found that people are influenced by superordinate spatial relations of geographic locations. They asked subjects to judge the spatial relations of different cities (e.g., Is Montreal north or south of Seattle?). People often misjudged the location of different cities based on the state or country the city is situated in (e.g., Montreal is in Canada and Canada is north of the United States, therefore Seattle must be south of Montreal). The superordinate category of country influenced judgements about the subordinate category of cities, irrespective of their spatial locations.

McNamara et al.⁶⁰ submitted subjects' recall protocols of a spatial layout to an algorithm that constructs hierarchical trees consistent with the internal organization of the protocols. They found that subjects' distance estimations and spatial priming was dependent on whether pairs of objects appeared in the same subtree or different subtrees of the hierarchy rather than each object's spatial location in the total array.

Both of the above studies highlight the importance of categorization and meaning in the organization of mental representations of spatial information. If hierarchical structure is important to the retrieval of spatial memories of objects in physical space, it is plausible that hi-

erarchical structure is even more relevant than spatial position to the recall of objects in hyperspace, where objects have no absolute position or even consistent relative position.

The development of spatial cognition

Much of the research on early spatial cognitive development focused on the development of children's cognition related to small spaces, particularly, the analysis of the children's changing frames of reference. As they develop, children switch from egocentric frames of reference to nonegocentric or geographic frames of reference.⁶¹ Egocentric references are spatial indicators pertinent to the child's own body and point of view (e.g., to the left or right of me). Geographic references are spatial indicators that are tied to the space as a whole (e.g., the relation between two objects in a room or the walls of a room). Children typically switch their frame of reference from egocentric to geographic references between 6 and 16 months of age. Beyond 16 months, changes in the choices of different geographic reference systems develop; that is, as they get older children increasingly choose more and more distal versus proximal systems (e.g., the walls of a room vs. the furniture in a room). As children's frames of reference radiate further away from their person, and include object-object relations as well as body-object relations, they are increasingly able to deal with larger environments.⁶² By 12 years of age, children's competence in real-world large-space navigation matches that of adults.⁵¹

Spatial knowledge acquisition in adults. In a highly influential article, Siegel and White⁶³ suggest that the ability to act efficiently within large-scale environments is based on the mental integration of multiple views. Because large environments cannot be viewed all at once, some type of mental-spatial representation must be constructed to comprehend the entire space. Siegel and White proposed that the acquisition of these spatial representations follows a distinctive course, as a series of stages, landmark recognition being the first in which individuals learn to discern and remember separate landmarks. In the second stage, land-

marks are coordinated in a sequence, which results in path or route representations. This route knowledge is comprised of both spatial and temporal relations between landmarks. As route representations become more refined, a representation of the entire space forms, a survey representation that includes the relations between landmarks and routes. Thus, with increased familiarity of the environment, spatial representations gradually incorporate more sophisticated versions of landmark, then route, then survey knowledge; survey knowledge being qualitatively different from the first two because it incorporates the understanding of metric distance relations between landmarks and routes. Siegel and White's view remained the dominant framework throughout the 1970s and 1980s, and has been cited in the hypertext literature.^{15,30,64}

More recently, Montello⁶⁵ argued that the final acquisition of survey knowledge is a quantitative rather than a qualitative change. He asserts that metric knowledge is being encoded right from the earliest learning stages, and that survey knowledge appears as the qualitative outcome of the integration of spatial knowledge, including metric knowledge, about separately learned places. Montello describes spatial knowledge as consisting of both nonmetric and metric information beginning with the first exposure to a new place. He also asserts that some types of spatial information are stored in a linguistic format and comes to the conclusion that "there exist two or more distinct subsystems that represent spatial information in different formats" (p. 17). Montello cites several authors who present converging evidence and arrive at comparable conclusions. For example, McNamara et al.⁶⁶ referred to three formats of spatial knowledge—temporal strings, propositions, and metric spatial representations. Kuipers⁶⁷ proposed that spatial knowledge is stored in many disconnected components that include separate metric and topological components. Kosslyn⁶⁸ suggested that the right hemisphere of the brain is specialized for processing coordinate spatial representations while the left hemisphere is specialized for processing categorical representations.

Three questions emerge from Montello's position: (1) If metric knowledge is so pervasive

throughout the acquisition of knowledge of large spaces, how relevant is real world navigation to hypertext as metric distance may not be a property of hypertext?; (2) Do spatial aids prompt the user to develop some form of mental representation that includes metric distances between elements?; (3) Integration of multiple types of spatial information is cognitively demanding⁶⁹; therefore, what is gained by adding components to hypertext (e.g., spatial maps) that may force another knowledge subsystem (in this case metric knowledge) to be engaged?

An alternative perspective is that navigation may be route-based, that is, route knowledge is not a prerequisite for or subset of survey knowledge but rather route and survey knowledge are acquired and represented separately. Hirtle and Hudson⁷⁰ found that subjects could acquire route knowledge from both slide presentations and maps but could only acquire survey knowledge from maps. This suggests that, with hypertext, even without global survey knowledge of a document, users should be able to encode sequences of landmarks into routes and integrate separate routes for a local representation of parts of the document. However, even with a route-based explanation of hypertext use, we still encounter the problem of distance definition. A route is defined by the sequence of and distances between landmarks. In hypertext, distance is not easily measured and is inconsistent from one traversal of the document to the next.

A PRELIMINARY COMPARISON OF HYPERSPACE AND PHYSICAL SPACE FEATURES

Metric distance may be one property that is not shared by hyperspace and physical space. In the following section several other possible disparities are suggested as examples of the types of issues that need to be investigated to evaluate the appropriateness of the spatial metaphor.

Perceptual cues

As mentioned previously, landmarks are important in navigation because they are a cue for a specific action to be carried out. In real space,

the surrounding environment provides additional retrieval cues for that action. In hypertext, we view only one window at a time. If the window itself is not a strong enough cue to retrieve the action from memory (e.g., select a particular icon to another window) and there are no navigational aids to help orient the user, the user must make a random selection of icons.

In contrast, real-world environments have an abundance of perceptual information; that is, they provide a rich collection of features and attributes on which we can base spatial decisions. Perceptual information allows us to make predictions about the future positions of objects in the environment as we move through space. For example, in a city setting, distal landmarks such as tall buildings are occluded by closer buildings. As we move through space, the degree of occlusion changes, giving us clues as to the future position of the buildings that are not entirely in our view in that moment.

Hypertext environments are featurally impoverished compared to real-world environments. This may make the attributes that are present more salient, but it also provides fewer opportunities for the user to establish multiple reference points that designate the user's position along a route. There is no perceptual information that is inherent to the hyperspace environment that helps us make predictions about the future positions of objects.

The importance of children developing broader frames of reference outside of their own bodies was mentioned earlier in the section on spatial cognition. The importance of using the body as a reference for spatial decision-making has been argued for adults as well.⁷¹ Body position as a perceptual cue and the sense-of-self within the environment as a physical cue are not characteristics of hyperspace.

The nature of landmarks

In real space, landmarks have no inherent meaning. Their meaning is defined by the action attached to them. Landmarks have a sequential relationship but not a semantic relationship. For example, a tree has no shared meaning with a car unless an action pairs the two together such as, "I must walk to the car and turn left toward the tree." In hypertext,

landmarks, such as individual windows, have their own semantic content that is related to other landmarks aside from their possible sequential or spatial relationship.

A sequence of landmarks makes up a route. In real space, the length of that route and the spaces between landmarks can be defined by a distance. Distance in real space is correlated with time, that is, the longer the distance the more time it takes to traverse it. In hyperspace, this correlation does not hold, partly because we have no way of accurately measuring distance but also because hypertext is dynamic. Distances do not remain static but change depending on the order of access the individual uses at any given time or the different preferences expressed by different users.

Spatial mental representation

Survey knowledge in a large hypertext document can only be achieved through navigational aids such as maps and content lists. This means that it is possible that survey knowledge is mentally represented in a different form than route knowledge where the user has visually experienced the pages. If survey knowledge is encoded in a different form than route knowledge, it must undergo some kind of mental manipulation so as to track it to route knowledge. This increases the cognitive demands of the task.

Contrary to this, in real world navigation distal landmarks and the general lay of the landscape are visible. Through these cues, survey knowledge can be partially acquired in the same manner as route knowledge; it is an expansion of the cognitive map developed for route knowledge. The accumulation of metric and nonmetric information occurs from the outset, that is, survey knowledge begins to be encoded in the same manner, at the same time as route knowledge. A presentation of a survey map in the homepage of a document provides survey information at the outset as well, but in a different form than route knowledge.

A question of scale

If physical space is cognitively related to hyperspace, how many nodes does a hypertext document need to have to approximate the

mental space equivalent to the size of a neighborhood or a city or province? In the real world, we have different cognitive strategies for dealing with different sizes of spatial configurations. How well hypertext cognitively maps onto physical space may be dependent on appropriate scale matches.

HYPertext NAVIGATION

To this point, this paper has focused on a broad overview of some of the issues regarding the spatial metaphor. The following section presents specific studies focused on the contribution of spatial navigational aids to hypertext performance.

Information access in hypertext documents is impacted by the structure of the document itself outside of the navigational aids that may be provided.⁷² There are three levels at which this may occur: (1) The inherent structure of the material (Is hierarchical organization an intrinsic quality of the material regardless of presentation?); (2) The organization of the material within the document or document structure cues (Are pages from subordinate categories differentiated from pages in superordinate categories?); and (3) The linking structure (Does it reflect both of the above or, in and of itself, does it impose structure or constraints on the document?). These are global features of the hypertext document that are beyond the focus of this paper. However, we need to be aware of and explicitly report these features during the course of investigating navigational aids.

Measures

How is navigational performance typically measured? Two categories of measurement encapsulate most of the measures typically used in hypertext research: efficiency and effectiveness measures.⁷³

Efficiency measures are based on speed and the number of steps taken to complete an information search. Efficiency measures often include the number of nodes visited in relation to the entire document, the number of excess nodes visited, the number of nodes revisited, the time taken on each node of the document, and the total task time.

Effectiveness measures focus on the user's search accuracy as well as his or her recall and understanding of the structure of the document. Effectiveness measures typically include measures of recall of the elements within and the structure of the document. This would include the number of node titles remembered, number of node titles filled into a blank overview, or a sketch map of the document.¹² Accuracy is also an effectiveness measure denoted by the number of correct answers to questions that require the user to conduct an information search.

A distinction needs to be made between navigation effectiveness and learning of the document material. In contrast to effectiveness, learning involves a deeper understanding of the content of the document and is often measured by the study or review time per node, the number of main ideas or concepts retained,⁷⁴ and the number of concepts retained over time.⁹

It is meaningful to clarify the distinctions between efficiency measures, effectiveness measures, and enduring learning as the first two are very important to information access in hypertext, while the third measure is pertinent to educational value.

Several studies that have found positive effects for spatially based navigational aids have reported an increase in efficiency measures but not in effectiveness measures.^{8,9,14} If hypertext

is to be a truly useful informational tool both measures would seem important. Making decisions quickly and accessing the least windows necessary is relevant, but if users are not able to find the information they need (incomplete or inaccurate search) or recall how they found it, the usefulness of hypertext could be severely constrained. Therefore, it would seem that a successful navigational aid should increase both efficiency and effectiveness.

Types of navigational aids

A spatial map, spatial overview, or graphical browser is a visual representation of the structure of the document that depicts the hypertext nodes and the links that connect them.^{8,9,74} These are usually in a diagrammatic form such as block diagrams, diagrams organized around a central term (spider map), or hierarchically ordered tree diagrams (Figs. 1 and 2).

Nonspatial navigational aids or textual aids include alphabetical or arbitrarily ordered content lists, conceptually ordered content lists, and indices.⁷⁵ It is difficult to pinpoint what features distinguish a textual from a spatial aid. I suggest that lists that assign a spatial location through features such as indentation⁷⁶ should be thought of as a spatial aid as they provide a spatial location for each term in the layout.

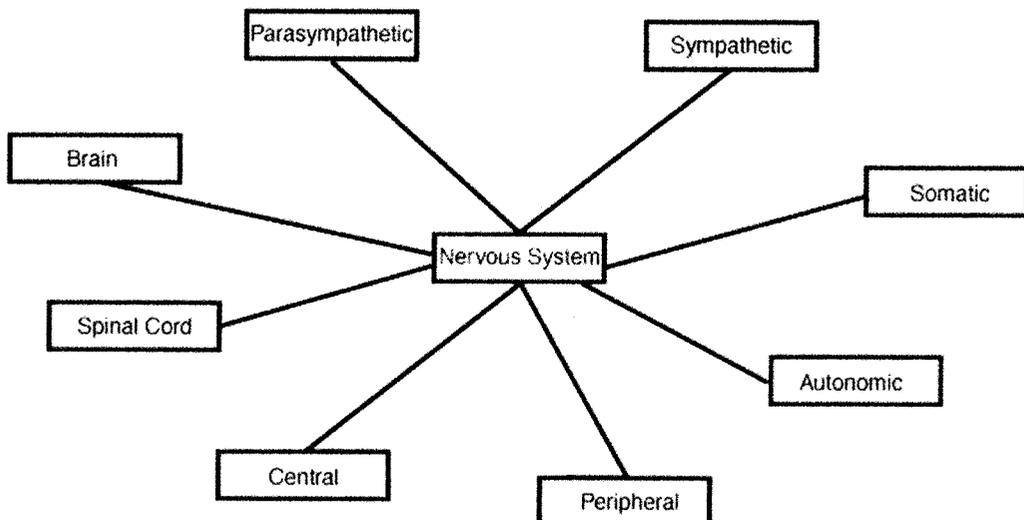


FIG. 1. An example of a spatially ordered navigation aid.

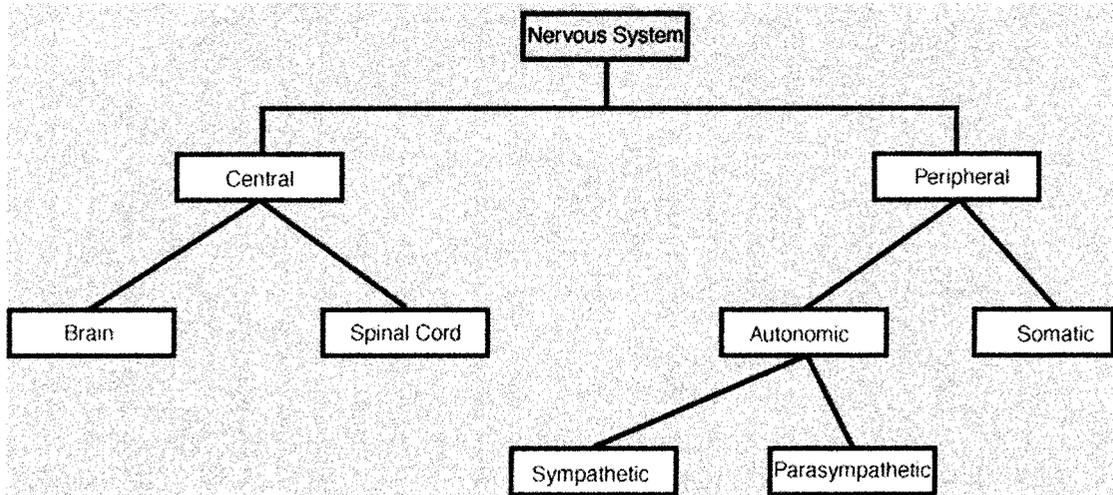


FIG. 2. An example of a spatial-hierarchically ordered navigation aid.

The effects of navigational aids

Most hypertext tasks can be described within the broad classification of open or closed tasks.⁷⁷ Open tasks have a general goal of browsing or exploring the document. Closed tasks have a specific goal of finding a particular piece of information. The effects of navigational tools on both types of tasks have been studied. The following section describes first a general metaanalysis on navigational aids, task types, and user characteristics, followed by specific examples of studies on closed and open tasks.

Chen and Rada⁷³ conducted a metaanalysis on 20 experimental studies and 3 doctoral dissertations that appeared between 1988 and 1993. Generally, the metaanalysis indicated that the effects of navigational tools such as indices, table of contents, and graphical maps interact strongly with the type of task involved (open or closed).

More specifically, for effectiveness scores (based on various measures of accuracy), task complexity showed the greatest effect size, followed by graphical maps. For efficiency scores (based on various speed measures), task complexity and spatial ability showed the largest effects. Also, spatial ability interacted with navigational aid on measures of efficiency (speed).

These results suggest that spatial processing may play a role in hypertext processing. However, as a metaanalysis is an analysis of common features across studies, it does not make

detailed comparisons of the specific features of each type of navigation tool. It is still not clear which specific features may contribute to this effect.

Effects of navigational aids on closed tasks

McDonald and Stevenson⁹ reported on two studies examining the effects of navigational aids on navigation and learning. In the first study, the effect of a localized spatial map, a textual content list, and no aid were compared. The map consisted of labels with connecting lines indicating the links between nodes. Not all nodes were linked. The content list was a scrollable list of all the nodes in the document. The navigational aids were not interactive in that users could not access pages via the aids but had to navigate through the pages indicated by the map or list. The aids were displayed in separate windows from the text content of the document. Subjects were asked to first find 10 target nodes and immediately after were asked 20 questions regarding information in the text. The navigation measures were: the mean time taken to access the target nodes and the number of additional nodes opened beyond that determined to be the optimal route. The learning measures were: the number of questions answered correctly and the mean number of node titles recalled during free recall. The results show that subjects in the map condition outperformed both the content list and no aid conditions on both measures of

navigation efficiency. However, no significant differences were found between the three conditions on accuracy on the test questions. The number of titles recalled were comparable between the spatial map and the contents list but were significantly less for the no aid condition. The spatial map did not produce better recall than the content list.

In McDonald and Stevenson's second study, a spatial map and conceptual map were compared to a no aid condition. The conceptual map was similar to the spatial map described for the first study but with the added feature of link descriptions. The same measures were recorded. Subjects again found 10 target nodes but then answered 40 questions based on the text, 20 factual questions and 20 questions meant to tap into deeper conceptual learning. This session was referred to as the acquisition phase. A week later, subjects returned to answer 40 additional questions to test for long-term retention (the retention phase). As in the first study, significant differences occurred between all conditions on both efficiency measures. On the effectiveness measures, there was no difference between the spatial and conceptual map conditions on the accuracy of factual questions in the acquisition phase. Both were better than the no aid condition. In the retention phase, conceptual map subjects answered more factual questions correctly than the spatial map and no aid condition. Regarding the conceptual questions, in the acquisition phase, subjects in the conceptual map and no aid conditions answered more questions correctly than the spatial map subjects. In the retention phase, there was a difference between all three groups; conceptual map subjects answering correctly most often, followed by the no aid group, and finally the spatial map group. These results suggest that a conceptual map is important for learning, whereas a spatial map increases efficiency but not effectiveness.

Dias and Sousa¹³ studied the effects on performance of a navigational map, described as a content index that provides a global overview of the document as well as a representation of the path taken by the user. Subjects were asked to find the answers to 12 questions. A map icon allowed the users to access the overview. From the overview, the subjects could access any

page within the document. Performance was measured on the accuracy of the information search (answers to the test questions), time spent on each screen page, frequency of visits to each information node, frequency of visits to the map, and time spent on the navigational map. Performance on these measures was compared to a predefined, optimal path for that particular test item as determined by the authors. Subjects were divided into three groups according to map access (both frequency and time spent on the map). Subjects who visited the map frequently and/or spent more time on the map accessed more unnecessary nodes than other users. The authors conclude that the "frequency of visiting the map and time spent on it do not guarantee orientation to the process of organizing and acquiring relevant information" (p. 184).

In another search task study, Leventhal et al.¹⁵ found that measures of navigational performance were not correlated with use of a hierarchical overview. Correlations between accuracy and speed measures and both the percentage and total number of visits to two types of hierarchical overview cards were not significant.

Stanton et al.¹² found evidence against facilitation of navigation maps on effectiveness measures. Subjects were divided into map and no-map conditions and were asked to search out the answers for a sentence completion task. They were then asked to draw an outline of the hypertext document on paper. The results indicated that the no-map subjects outperformed the map subjects on accuracy of the sentence completion task. Subjects in the no-map condition were also better able to produce a drawing of the document structure and reported feeling more in control of the search situation.

Effects of navigational aids on open tasks

Dee-Lucas and Larkin⁷⁴ investigated the effects of structured (a map) versus unstructured (an alphabetical list) overviews on recall. The map consisted of a hierarchical diagram with one central heading and three levels of sub-headings. Through the placement of the headings, referred to as unit titles, the map indicated the subordinate and superordinate relations

between the unit titles. The list provided no indication of the relations between the unit titles. Both overviews were interactive, in that as the reader finished a unit he or she was required to return to the overview to select and link up to a new unit title. Both document formats contained nine units. Participants were asked to read all material first in a specified order and then were free to review any material they chose. Participants were then tested on both recall and study strategy measures. The recall measures included the number of unit titles recalled, the proportion of text propositions (main ideas) recalled, the proportion of propositions recalled per unit, and the proportion of units with at least one proposition recalled. On these measures of recall there were no differences between the map and the list groups. However, the map group were more able to recall the title locations than the list group, and the list group spent more time on the overview.

Wenger and Payne¹⁴ investigated whether the presence of a graphical browser impacted the comprehension and retention of the content and structure of a hypertext document. In this study, the graphical browser indicated the structure of the document in a hierarchical fashion and also, through the highlighting of visited nodes, allowed the user to see which portions of the document had been accessed. To minimize imposing a certain order of viewing, users were allowed to access the document at any entry point by choosing, on the first node, from an unstructured alphabetical listing of the contents. Both the browser group and the no-browser group had this option. The browser group could return to the browser from any node in the document. The subjects were instructed to read through all the nodes and indicate when they felt they had accessed all available nodes. At the end of the text, in a timed recall task, users were required to note as many as the node titles as they could remember. They then completed multiple choice tests as an indicator of their comprehension of the text. Finally, to tap into the users' knowledge of the text structure, users were presented with pairs of node titles and asked to state whether the nodes were linked within the previous text. On all three of these effectiveness measures, results showed no differences be-

tween the browser and no-browser groups. The only significant differences occurred in efficiency measures: the number of nodes accessed and the number of nodes revisited. Subjects in the browser groups accessed a higher proportion of the total available nodes and exhibited less revisitation than the no-browser subjects. Subsequently, the authors suggested that presentation of a graphical browser enhanced the efficiency of browsing but did not improve users' performance on measures of effectiveness.

An important aspect to note about a number of the studies comparing spatial navigation tools versus content lists is that hierarchical organization and spatial location are often confounded.^{9,10,74} By placing node titles within a hierarchical framework, nodes are automatically assigned a spatial location. Conversely, providing a basic content list removes not only spatial location but also the hierarchical information that highlights the conceptual relations between node titles. Although any time one label is visually differentiated from another some sort of spatial context is created, this effect might be reduced by using typographical cues (e.g., colored or bolded font) on simple content list labels to denote which topics are nested within other topics (providing information about conceptual relations) while keeping at a minimum the assignment of an explicit spatial location. However, there are drawbacks to this method as well. The hierarchy of the document could be quite broad (have many labels at the same level of the hierarchy), but could not be very deep (have many levels within the hierarchy) or the list itself would become overly complex.

THE ROLE OF COGNITIVE PSYCHOLOGY

The previous section provided examples of variables that can contribute to navigation outcomes. Task type, navigational aids, and overall document structure all impact information search. The section also illustrates through empirical studies the inconclusiveness of data toward answering the central questions raised in this article.

The following section is an outline of some ways in which traditional cognitive psychology paradigms could be used to investigate some of the issues described throughout this paper.

Spatial mental representations: Accessing underlying mental representations and matching to the conceptual properties of the material

Do spatially based aids facilitate the construction of a mental representation of the document that reflects the organization of the material? The same methodologies used to examine spatial cognition in physical space could be applied to examine cognitions in hyperspace. For example, McNamara et al.'s⁶⁰ hierarchical tree analysis for assessing the underlying spatial representations for physical spatial layouts could be applied to users' mental representations of hypertext documents. Subjects' recall protocols of the hypertext document structure could be subjected to McNamara et al.'s algorithm, which constructs hierarchical trees consistent with the internal organization of the protocols. Subjects' distance estimations and spatial priming could be used to determine whether pairs of windows that are conceptually related appeared in the same subtree or different subtrees of the hierarchy, rather than in the order of access or some other type of spatial organization. When comparing navigation aids with different degrees of spatial information, this method could reveal which is most important to the user, conceptual or spatial information.

The role of spatial processing: Identifying the presence of spatial processing

Is spatial processing involved in the execution of a hypertext search task? The dual-task interference paradigm is often used in cognitive psychology to detect the presence of specific types of cognitive processing.⁷⁸ The basis of this task is that two tasks executed concurrently will produce degraded performance, particularly if they involve the same type of processing. When a subject is presented with a second task that requires a different type of processing than the task already engaged in, the second task will interfere less with performance than a task that demands the same type

of processing. Particular to hypertext, if subjects are asked to engage in a hypertext search task and this involves spatial processing, then a second task also involving spatial processing should degrade performance on the search task outcomes. Conversely, a second task involving some other type of processing (e.g., auditory processing) should not impact performance to such a degree. If information in hypertext is encoded spatially, the spatial task will interfere with the encoding of the document structure and affect measures such as the recall of node titles. The spatial task could involve monitoring the appearance and location of a visual object on screen. The nonspatial task could involve distinguishing several auditory tones from one another.

Metaphor validation: Mapping features of physical space to features of hyperspace

Is the spatial metaphor appropriate for hypertext? It is prudent to take a closer look at the benefits and limitations of this all-encompassing spatial metaphor from the same perspectives (operational, structural, and pragmatic) and with the same rigor that are used to evaluate metaphors of a narrower focus. Identifying the mismatches is important. Even if users may not be able to report them, mismatches may provide implicit knowledge to users that may negatively affect their mental representations of the hypertext environment.

Initially, a structural analysis could be undertaken, as this is the type of metaphor analysis often used in both psychological and HCI research. Also, structural feature mapping between hyperspace and physical space would be useful toward setting the parameters of the metaphor, a framework on which to pursue the operational and pragmatic analyses. A question raised earlier may be a relevant starting point: Is metric distance a property of hypertext? If metric distance is a property of hypertext then the knowledge we have about how people navigate in physical space may be relevant to navigation in hyperspace. One way of investigating this issue is to determine whether hyperspace violates the three axioms that define physical space. For example, the symmetry axiom could be evaluated through the col-

lection of distance ratings from subjects asked to make comparisons between pairs of windows. Distance ratings for all pairs and pair orders (e.g., the distance from A to B as well as the distance from B to A) could be collected and the matrix of this data could be evaluated for symmetry.

If the data does not support a metric representation (e.g., shows asymmetric distances), further evaluation based on adherence to the spatial primitives outlined by Golledge⁴⁶ would be revealing. Golledge is cautious about the application of spatial primitives as he defines them for several reasons. Some of these definitions can only be expressed and investigated using imprecise, natural language (e.g., How can we quantify identity?) rather than in the unambiguous technical terms used in computer spatial systems such as Global Positioning Systems. Golledge expresses concern that such systems entail technical language based on imprecise definitions of spatial primitives:

We do this without really knowing how error accumulates as spatial primitives are joined or as spatial derivatives are deduced from spatial primitives. We do this without fully understanding that our fuzzy and error-ridden concepts might be inappropriately transferred to the exact and precise domain of the computer. (p. 103)

Therefore, examining the comparison between hyperspace features and spatial primitives may be best undertaken as a combination of quantitative and qualitative analyses. Data collected via tracking mechanisms would include quantitative variables such as the time spent on individual windows, the number of times a window is revisited, and the number of excess windows opened beyond an optimal path. To interpret users' behaviors, quantitative data such as the variables above may be combined with data from spatial products. For example, frequent revisitations to a window may indicate that a particular window is being used as a reference point or landmark. This implies that a window is being assigned a location relative to other windows. If a window is used as a landmark, sketch maps of networks may portray that window as having a central

location with other windows branching off. An alternate view is that revisitation may be a symptom of confused identity. Because of poorly differentiated labels or visual appearance, a subject may revisit a window mistaking it for a new unopened window or confusing it with a different, previously accessed window. Again, including qualitative data may be illuminating. Verbal justifications of revisitations could be collected during the search tasks to tap into these differences.

Also, verbal protocol recording can be undertaken during the navigation task. Subjects could be trained to "think aloud" their strategies and thoughts as they proceed through a hypertext document. Previous research on protocol analysis suggests the task of navigating a document would be conducive to the "think aloud" approach and produce the type of protocol data that can be analyzed in detail.⁷⁹

CONCLUSION

From the above review it is apparent that there are many complex psychological factors involved in the use of the World Wide Web. Because of the vast potential for the Web to impact millions of people, it is important that we develop a detailed understanding of users' perceptions of hyperspace, beginning with a thorough investigation of the most basic cognitive processes involved, processes that have been and continue to be closely examined with respect to many other types of environments and stimuli.

For the computer scientist's goal of developing optimal interfaces, information gathered from such cognitive studies should aid in the development of more user-friendly web sites, navigation tools, and search mechanisms. For clinical psychologists, such information may highlight aspects that need to be considered in the Web presentation of psychological services. This information may also shed light on the inherent aspects of hyperspace that allow for the development of negative behaviors associated with the Web. For social psychologists, understanding hyperspace and how it cognitively maps onto physical space may predict behavioral advantages, disadvantages, and solutions

related to issues about the globalization of our society and the individual's social place within it. For educators, such information may facilitate the presentation of instructional material in a format that will reduce cognitive overhead in the learner.

Understanding the cognitive processes involved in hypertext use, particularly people's basic conceptualizations of hyperspace, can only enhance our ability to make the World Wide Web an accessible information resource for users of all types. Cognitive psychologists can play an important role in the identification and analysis of how the processes of the mind interact with the Web.

REFERENCES

1. Rock, I. (1983). *The logic of perception*. Cambridge, MA: MIT Press.
2. Goodman, N. (1978). *Ways of worldmaking*. Indianapolis, IN: Hackett Publishing Co.
3. Holyoak, K.J., Thagard, P. (1996). *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press.
4. Pylyshyn, Z.W. (1993). *Metaphorical imprecision and the "top-down" research strategy*. In: Ortony, A., (ed.) *Metaphor and thought*. Cambridge, MA: Cambridge University Press, pp. 420-436.
5. Canter, D., Rivers, R., & Storrs, G. (1985). Characterizing user navigation through complex data structures. *Behavior and Information Technology*, 45(2):93-102.
6. Kim, H., & Hirtle, S.C. (1995). Spatial metaphors and disorientation in hypertext browsing. *Behavior and Information Technology*, 14(4):239-250.
7. Kim, J. (1999). An empirical study of navigational aids in customer interfaces. *Behavior and Information Technology*, 18:213-224.
8. McDonald, S., & Stevenson, R. (1998). Spatial versus conceptual maps as learning tools in hypertext. *Journal of Educational Hypermedia and Multimedia*, 8(1): 43-64.
9. McDonald, S., & Stevenson, R. (1999). Navigation in hyperspace: An evaluation of the effects of navigational tools and subject matter expertise on browsing and information retrieval in hypertext. *Interacting with Computers*, 10:129-142.
10. Beasley, R.E., & Waugh, M.L. (1995). Cognitive mapping architectures and hypermedia disorientation: An empirical study. *Journal of Multimedia and Hypermedia*, 4:239-255.
11. Schroeder, E.E., & Grabowski, B.L. (1995). Patterns of exploration and learning with hypermedia. *Journal of Educational Computing Research*, 13(4):313-335.
12. Stanton, N.A., Taylor, R.G., & Tweedie, L.A. (1992). Maps as navigational aids in hypertext environments: An empirical evaluation. *Journal of Educational Multimedia and Hypermedia*, 1(4):431-444.
13. Dias, P., & Souza, A.P. (1997). Understanding navigation and disorientation in hypermedia learning environments. *Journal of Multimedia and Hypermedia*, 6(2):173-185.
14. Wenger, M.J., & Payne, D.G. (1994). Effects of a graphical browser on readers' efficiency in reading hypertext. *Technical Communications*, 41:224-233.
15. Leventhal, L.M., Teasley, B.M., Instone, K., Rohlman, D.S., & Farhat, J. (1993). Sleuthing in Hyperholmes-super™: An evaluation of using hypertext vs. a book to answer questions. *Behavior & Information Technology*, 12(3):149-164.
16. Reynolds, S.B., & Dansereau, D.F. (1990). The knowledge hypermap: An alternative to hypertext. *Computers in Education*, 14(5):409-416.
17. Campagnoni, F.R., & Ehrlich, K. (1990). Getting there when you don't know where "there" is: Navigational strategies in a hypertext help system. *SIGCHI Bulletin*, 21(3):17-18.
18. Vicentes, K.J., & Willeges, R.C. (1988). Accommodating individual differences in searching a hierarchical file system. *International Journal of Man-machine Studies*, 29:647-668.
19. Leidig, P.M. (1992). The relationship between cognitive styles and mental maps in hypertext assisted learning. *Dissertation Abstracts International*, 53(5-A): 1372.
20. Bieber, M., Vitali, F., Ashman, H., Balasubramanian, V., & Oinas-Kukkonen, H. (1997). Fourth generation hypermedia: Some missing links for the World Wide Web. *International Journal of Human-Computer Studies*, 47:31-65.
21. Jonassen, D.H. (1986). Hypertext principles for text and courseware design. *Educational Psychologist*, 214: 269-292.
22. Jonassen, D.H. (1989). *Hypertext/hypermedia*. Englewood Cliffs, NJ: Educational Technology Publications.
23. Schneiderman, B., & Kearsley, G. (1989). *Hypertext hands-on: An introduction to a new way of organizing and accessing information*. NY: Addison-Wesley.
24. Conklin, J. (1987). Hypertext: An introduction and survey. *Computer, IEEE Computing Society*, 20(9):17-41.
25. Nielsen, J. (1990). *Hypertext and hypermedia*. San Diego, CA: Academic Press.
26. Gygi, K. (1990). Recognizing the symptoms of hypertext . . . and what to do about it. In: Laurel, B., (ed.) *The art of human-computer interface design* New York: Addison-Wesley, Reading, pp. 279-288.
27. Tsai, C.L. (1988). Hypertext: Technology, applications and research issues. *Journal of Education Systems* 17(1):3-14.
28. Ransom, S., Wu, X., & Schmidt, H. (1997). Disorientation and cognitive overhead in hypertext systems. *International Journal of Artificial Intelligence Tools*, 6(2):227-253.

29. McLeese, R.A. (1989). Navigation and browsing in hypertext. In: McLeese, R.A., (ed.) *Hypertext: Theory into practice*. Westport, CT: Ablex, pp. 6–44.
30. Edwards, D.W., & Hardman, L. (1989). "Lost in hyperspace": Cognitive mapping navigation in a hypertext environment. In: McLeese, R.A., (ed.) *Hypertext: Theory into Practice*. Westport, CT: Ablex, pp. 90–105.
31. Foss, C.L. (1989). Detecting user's lost: Empirical studies on the browsing of hypertext. *Technical Report*, No. 972, INRIA, Sophia-Antipolis.
32. Hammond, N., & Allinson, L. (1987). The travel metaphor as design principle and training aid for navigating around complex systems. In: Diaper, D., Winder, R., (eds.) *People and Computers III: Proceedings of the third conference of the British Computer Society, Human-Computer Interaction Specialist Group*. Cambridge: Cambridge University Press.
33. Shum, S. (1990). Real and virtual spaces: Mapping from spatial cognition to hypertext. *Hypermedia*, 2(2): 133–158.
34. Jones, W.P., & Dumais, S. (1986). The spatial metaphor for user interfaces: Experimental test of reference by location versus naming. *ACM Transactions on Office Information Systems*, 4:42–63.
35. Stanton, N.A., & Baber, C. (1994). The myth of navigating in hypertext: How a "bandwagon" lost its course! *Journal of Educational Multimedia and Hypermedia*, 3:235–249.
36. Mayes, T., Kibby, M., & Anderson, T. (1990). Signposts for conceptual orientation: Some requirements for learning from hypertext. In: McLeese, R., & Green, C., (eds.) *Hypertext: State of the art*. Oxford: Intellect.
37. Lakoff, G. (1994). What is metaphor? In: Barnden, J.A., Holyoak, K.J., (eds.) *Advances in connectionist and neural computation theory: Analogy, metaphor and reminding*. New Jersey: Ablex.
38. Carroll, J.M., Mack, R.L., & Kellogg, W.A. (1988). Interface metaphors and interface design. In: Helander, M., (ed.) *Handbook of human-computer interaction*. North-Holland: Elsevier Science Publishers, pp. 67–85.
39. Blumenthal, L.M. (1953). *Theory and applications of distance geometry*. London: Oxford.
40. Shepard, R.N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. Part I. *Psychometrika*, 27:219–246.
41. Kruskal, J.B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29:115–129.
42. Guttman, L. (1971). Measurement as structural theory. *Psychometrika*, 36:465–506.
43. Carroll, J.D., & Wish, M. (1974). Multidimensional perceptual models and measurement methods. In: Carterette, E.C., Friedman, M.P., (eds.) *Handbook of perception*, Vol. 2. New York: Academic Press, pp. 391–447.
44. Tversky, A. (1977). Features of similarity. *Psychological Review*, 84:327–352.
45. Tversky, A., & Gati, I. (1982). Similarity, separability and the triangle inequality. *Psychological Review*, 89:123–154.
46. Golledge, R.G. (1995). Primitives of spatial knowledge. In: Nyerges, T., et al. (ed.) *Cognitive aspects of human-computer interaction for geographic information systems*. NATO/ASI Series: Kluwer, pp. 29–44.
47. Passini, R. (1984). *Wayfinding in architecture*. New York: Van Nostrum Reinhold.
48. Cornell, E.H., & Heth, C.D. (1983). Spatial cognition: Gather studies used by preschool children. *Journal of Experimental Child Psychology*, 35:93–110.
49. Shepard, R.N., & Metzler, J. (1971). The mental rotation of three-dimensional objects. *Science*, 17:701–703.
50. Acredolo, L.P. (1978). Development of spatial orientation in infancy. *Developmental Psychology*, 14:224–234.
51. Cornell, E.H., Heth, C.D., & Alberts, D. (1994). Place recognition and wayfinding by children and adults. *Memory and Cognition*, 22(6):633–643.
52. Liben, L. (1981). Spatial representation and behavior: Multiple perspectives. In: Liben, L.S., Patterson, A.M., & Newcombe, N. (eds.) *Spatial representation and behavior across the life span*. New York: Academic Press, pp. 16–37.
53. Tolman, E.C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55:189–208.
54. Presson, C.C., & Somerville, S.C. (1985). Beyond egocentrism: A new look at the beginnings of spatial representation. In: Wellman, H.M. (ed.) *Children's searching: The development of search skill and spatial representation*. Hillsdale, NJ: Erlbaum, pp. 1–26.
55. Lieblisch, I., & Arbib, M. (1982). Multiple representation of space underlying behaviour and associated commentaries. *The Behaviour and Brain Sciences*, 5(4):627–660.
56. Kaplan, S. (1973). Cognitive maps in perception and thought. In: Downs, R.M., Stea, D., (eds.) *Image and environment*. Chicago, IL: Aldine, pp. 63–78.
57. Golledge, R.G., & Stimson, R.J. (1997). *Spatial behavior: A geographic perspective*. New York: Guilford Press.
58. Kitchin, R.M. (1994). Cognitive maps: What are they and why study them? *Journal of Environmental Psychology*, 14:1–19.
59. Stevens, A., & Coupe, P. (1978). Distortion in judging spatial relations. *Cognitive Psychology*, 10:422–437.
60. McNamara, T.P., Hardy, J.K., & Hirtle, S.C. (1989). Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15:211–227.
61. Acredolo, L.P. (1985). Coordinating perspectives on infant spatial orientation. In: Cohen, R., (ed.) *The development of spatial cognition*. New Jersey: Erlbaum.
62. Cohen, R. (1985). *The development of spatial cognition*. New Jersey: Erlbaum, pp. 115–133.
63. Siegel, A.W., & White, S.H. (1975). The development of spatial representation of large-scale environments. In: Reese, H.W. (ed.) *Advances in child development and behavior*, Vol. 10. New York: Academic, pp. 9–55.
64. Dillon, A., McKnight, C., & Richardson, J. (1990) Nav-

- igation in hypertext: A critical review of the concept. In: Diaper, D., et al., (eds.) *Human-Computer Interaction—INTERACT'90*. North-Holland: Elsevier Science Publishers, pp. 587–592.
65. Montello, D.R. (1998). A new framework for understanding the acquisition of spatial knowledge in large-scale environments. In: Golledge, R.G., Egenhofer, M.J. (eds.) *Spatial and temporal reasoning in geographic information systems*. New York: Oxford University Press, pp. 143–154.
 66. McNamara, T.P., Halpin, J.A., & Hardy, J.K. (1992). Spatial and temporal contributions to the structure of spatial memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18:555–564.
 67. Kuipers, B. (1982). The “map in the head” metaphor. *Environment and Behavior*, 14:202–220.
 68. Kosslyn, S.M. (1987). Seeing and imagining in the cerebral hemispheres: A computational approach. *Psychological Review*, 94:148–175.
 69. Golledge, R.G., Ruggles, A.J., Pelligrino, J.W., & Gale, N.D. (1993). Integrating route knowledge in an unfamiliar neighborhood: Along and across route experiments. *Journal of Environmental Psychology*, 13:293–307.
 70. Hirtle, S.C., & Hudson, J. (1991). Acquisition of spatial knowledge for routes. *Journal of Environmental Psychology*, 11:335–345.
 71. Garling, T., Book, A., & Lindberg, E. (1989). Adults' memory representations of the spatial properties of their everyday physical environment. In: Cohen, R., (ed.) *The development of spatial cognition*. New Jersey: Erlbaum, pp. 141–177.
 72. Korthauer, R.D., & Koubek, R.J. (1994). An empirical evaluation of knowledge, cognitive style, and structure upon the performance of a hypertext task. *International Journal of Human-Computer Interaction*, 6(4): 373–390.
 73. Chen, C., & Rada, R. (1996). Interacting with hypertext: A meta-analysis of experimental studies. *Human-Computer Interaction*, 11:125–156.
 74. Dee-Lucas, D., & Larkin, J.H. (1995). Learning from electronic texts: Effects of interactive overviews for information access. *Cognition and Instruction*, 13(3):431–468.
 75. Simpson, A., & McKnight, C. (1990). Navigation in hypertext: Structural cues and mental maps. In: McLeese, R., & Green, C. (eds.) *Hypertext: State of the Art*. Oxford: Intellect.
 76. Chou, C., & Lin, H. (1997). *Navigation maps in a computer-networked hypertext learning system*. Paper presented at the annual meeting of the Association for Educational Communications & Technology, Albuquerque, New Mexico.
 77. Marchionini, G. (1989). Information-seeking strategies of novices using a full-text electronic encyclopedia. *Journal of the American Society for Information Science*, 40:54–66.
 78. Gopher, D. (1994). Dual-task performance: Rationale and scope. In: Eysenck, M.W., (ed.) *The Blackwell dictionary of cognitive psychology*. Cambridge, MA: Blackwell, pp. 111–116.
 79. Ericsson, K.A., & Simon, H.A. (1984). *Protocol analysis: Verbal reports as data*. Cambridge, MA: MIT Press.

Address reprint requests to:

Patricia M. Boechler
 Department of Psychology
 P108 Biological Sciences Bldg.
 University of Alberta
 Edmonton, Alberta, Canada T6G 2E9

E-mail: pmb2@ualberta.ca