

Learning From the City: Toward a Better Sense of Place in Information Space

by

Joshua Goldberg

B.S., Massachusetts Institute of Technology (1996)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1997

© Joshua Goldberg, 1997

The author hereby grants to MIT permission to reproduce and
to distribute copies of this thesis document in whole or in part.

Signature of Author
Department of Electrical Engineering and Computer Science
August 8, 1997

Certified by
Randall Davis
Associate Director, Artificial Intelligence Laboratory
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Departmental Committee on Graduate Theses

Learning From the City: Toward a Better Sense of Place in Information Space

by

Joshua Goldberg

Submitted to the
Department of Electrical Engineering and Computer Science
on August 8, 1997, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

The sense of the space and the sense of place are introduced as concepts to aid the design of information spaces. The sense of the space is a user's perception of how an information space is organized. The sense of place is the way a user perceives his or her own place in relation to the structure of the space. Ideally a user will have a comprehensive sense of the space and maintain only a minimal, local sense of place.

Four techniques are presented to help users to form a clear sense of the space. These techniques are demonstrated in an information space with a zooming and panning interface.

Additionally, three interface features are discussed that encourage users to maintain a simple sense of place. These interface features are incorporated into a World Wide Web site. An experiment is conducted examining user behavior in the site and demonstrating that users do indeed have a simpler sense of place in the new site.

Thesis Supervisor: Randall Davis

Title: Associate Director, Artificial Intelligence Laboratory

Acknowledgments

I would like to acknowledge the support of the following people:

First and foremost, my advisor Randall Davis who taught me countless lessons over the last year.

Whitman Richards, Stephen Gilbert, Brygg Ullmer, Ed Barrett, Tom, and Rosalita for many helpful suggestions. Mr. Death, my Netscape historian. Cecily, who was going to say, "It's a rock."

Max and Tom for showing interest. BDO for tolerating my prolonged absence. My typists, especially Tom who slept less than me, and the D.S.O. for paying most of them. (Please watch those hands, kiddies.)

Sol, Rose, Push, Janet, Mike, Jess, Tom, Tim, Tom, Bayard, and Eddie for much friendship. Mom, Dad, Stu and Dina for much love.

Erica for dragging me through the mud, after her bedtime no less. And rif, who started it all with his "Back" button.

This research was supported in part by the Defense Advanced Research Projects Agency and was monitored by the Office of Naval Research under contract number N00014-94-1-1090.

This thesis is dedicated to the memory of my teacher and friend Mr. Mac.

Chapter 1

Introduction

In an information space, navigational concerns are always competing for a user's attention with such activities as reading and comprehension, learning, or making informational comparisons. To decrease the impact of navigation on the capacities of users of information spaces, I hope to show how to create information spaces that are “navigationally obvious,”¹ supporting users to make navigational decisions immediately with only local knowledge and no previous navigational information. At a decision point it should be obvious what to do based only on what is locally visible and knowledge about the space. This benefits users by freeing up working memory for other activities between navigational decisions.

Hypertext browsers give users an impoverished notion of a navigational history. As users, all we can do is stack and unstack pages. Places we have visited are behind us and behind the places we have been since. Going back is always a regression and a determination that what we left behind is more important than where we are. We use the “Back” button like an Undo button to correct our navigational mistakes and restore the state. The limitations of the stack are significant because browsing history is not irrelevant to users. I argue in Chapter 3 that users take on the browser's history stack as a means of understanding their place. This goes counter to the goal of a minimal, easy-to-maintain sense of place; users must make adjustments to their conception of their place with each step.

¹David McAllester posed a formal notion of obviousness in his paper, “Observations on Cognitive Judgments.” [McA91] For example, it is obvious that in chess a king can get to any position on the chess board. It is not obvious that a knight can do the same. Informally, obvious statements are those that can be verified by making a series of local inferences, independent of one another.

To support a minimal sense of place, a user browsing a site needs comprehensive knowledge of the structure of the space and how to recognize places within it. The sense of the space, the topic of Chapter 2, is a user's conception of the structure and content of an information space. Hypertext users often keep track of the route they are following in order to be able to get back out the same way. This suggests a fear of venturing through uncharted territory. A good sense of the space lets users navigate so confidently that they may freely "lose themselves" in the space without fear of actually getting lost.

Since the space is static, users need not expend much effort while browsing to maintain a comprehensive sense of the space. A user's place, however, and therefore his or her sense of place, is always changing. Maintaining a comprehensive sense of place requires continual attention while browsing to keep it up to date. It is a distraction from whatever activity the space was meant for.

Chapter 2 presents four techniques I developed to enhance the comprehensibility and memorability of information space structure. Distinctive character makes parts of spaces more recognizable and memorable. Common character helps users to group objects together. Reminding vistas and countable character serve to teach and remind users how parts of a space fit into a global structure. I demonstrated these techniques using a space with a zooming and panning interface.

The sense of place is a user's working knowledge of where he or she is in a site. Chapter 3 presents three interface characteristics that engender a simpler sense of place in users. First, a largely static interface lessens users' feelings of immersion in the space, decreasing their worry of getting lost and thus decreasing their tendency to keep track of navigational history. Second, a continually visible place indicator lets users rely on the interface to maintain their sense of place. Finally, immediate top-level navigation prevents users from keeping constant track of how to get to a central home node by infusing the central navigational importance of home into the interface of every page.

Based on the assumption that the frequent use of the "Back" button indicates users are aware of their browsing history as they navigate, I showed experimentally that a site

designed with these three characteristics encouraged users to navigate more locally and immediately without so much of an eye to their navigational history.

Throughout this thesis I use the terms site, information space, and space interchangeably. These terms refer both to a collection of information and its interface. It is assumed that the collection is too large to fit on a manageably sized display. This means that they must be experienced temporally, bit by bit, by the users. The scale of the examples I used in this thesis is that of an organization's Web site. In particular, examples draw on web-based information about a set of projects at the MIT Artificial Intelligence Laboratory.

Chapter 2

A Sense of the Space

The sense of a space is the way a user visualizes the structure of a site based on personal experiences browsing the site. It consists of knowledge about the relationships between locations in the space, as well as cues for recognizing places. If a space does not clearly communicate its structure in a comprehensible and memorable way users may never acquire a coherent sense of the space. A designer's efforts to organize a collection of information will be wasted if users do not perceive that organization as they browse.

The benefits of a clear mental image of a site are many. Users will be better able to remember the contents of a site. They will know better where they want to go and how to get there. Memory of the structure of the site gives users a structure with which to remember the contents of the site.

In addition, a clear sense of the space allows a user to have a simpler sense of his or her place within the site. (The sense of place is explored in depth in chapter 3.) By understanding how the parts of the site fit together and by being able to recognize places in the site, users will have less need to remember the details of how they navigated from one place to another. A strong sense of space enables a user to navigate by simple, local steps where recognition triggers the next action. They will know where they are going and be less likely to feel lost.

The fruits of this chapter are four design techniques for making information spaces more structurally comprehensible and memorable: distinctive character, common character in conjunction with locality or path continuity, reminding vistas, and countable character.

These techniques are not organizational techniques to structure a space. They are rather ways to clarify the site to make its structure more noticeable, comprehensible, and memorable. At the end of this chapter I discuss how these four techniques were applied in the creation of a demonstration space using a zooming and panning interface.

2.1 Mental representations of physical spaces

This section reviews the major theories of how people perceive physical environments. The concept of a “map in the head,” a two-dimensional map-like mental representation of the world used for human spatial navigation, came into popularity among cognitive scientists after the 1948 publication of Edward Tolman’s influential report on the subject. Tolman definitively demonstrated using rats in mazes that simple stimulus-response behaviors could not account for the navigational learning evident even in rats. The “map in the head” was the theory Tolman proposed to replace stimulus-response as the means of navigation for “mice and men.” [Tol48]

It is now generally accepted that people do not recall spatial information by reading a mental map as Tolman had imagined. The memory required to recall a map would be too much of a burden at every step of navigation. [CC81] A map representation would never allow for the common phenomenon of a route learned in only one direction. [Kui82] Finally, if people understood spaces by mental maps, they would not make a well-known class of errors known as hierarchicalization errors. [CC81] [Kui82] An example of a hierarchicalization error is that many Americans have difficulty believing that Los Angeles is actually east of Las Vegas. People are adamant that that is not the case because they know that California is generally west of Nevada. In addition, remembering two dimensional structure to make each navigational decision to make each navigational decision would be a great burden. Tables associating the recognition of a place with facts relating that place to others has been proposed as an alternative to maps. Such tables would support route and survey knowledge using less memory. [Kui82]

Route knowledge and survey are the spatial representations most generally accepted by cognitive scientists. Route knowledge gives a sequence of actions to get from one place to another along a path. Survey knowledge organizes regions of the environment by synthesizing collections of spatial relationships. An example of a piece of survey knowledge is, “Kenmore Square in Boston is to the east of the Back Bay.” Route knowledge is used in automatic navigation. At each decision point, visual cues trigger long term memory recall of what to do next. Survey knowledge can be used to discover new paths through a less familiar environment. [CC81] [Tve93] [Kui82] [Kui78] A study of taxi drivers showed that as experts gain experience they learn more route knowledge and rely less on survey knowledge. [Cha83]

Although it should not be taken as a theory of representation, urban planners have found it useful to study a kind of two dimensional image, more abstract than Tolman’s “map in the head.” The urban planner Kevin Lynch’s seminal book *The Image of the City* [Lyn60] is a study of city dwellers’ perceptions of urban forms based on interviews and mapping exercises. Lynch showed using examples from studies of city dwellers how confusion about the environment can occur when these forms are not clear in the actual physical environment. Landmarks off-center in districts, edges that could not be followed for their entire lengths, indistinct districts, and a variety of other urban phenomena caused errors, distortions, or unrecalable voids in the maps constructed by inhabitants. On the other hand, features like vivid contrast and singularity, continuity, meaningful names and symbols, and vistas which increase the scope of vision help people to understand and remember the features of the environment and how they relate. Lynch’s discussion of these urban problems and assets was an explicit source of inspiration by analogy for my own development of techniques to create a comprehensible sense of the space in a virtual information space.

It is important to keep in mind that although Lynch’s work as well as further studies in the same vein such as [App69] is a valuable source of design guidance to route out environmental confusion, it should not be taken as a theory of how the mind represents spatial environments. First, because the results of the studies are maps that average the

information as recalled by a collection of subjects, the theory expresses only the major parts of the environment shared by everyone, [CC81] not, for instance, an individual's favorite convenience store with good bagels on the way to work. Further, a map drawn by an individual should not be interpreted as a mental representation. These maps are two-dimensional reconstructions based on information about the relationships between places that is not represented two-dimensionally. [CC81] [Kui82]

2.2 Techniques

In this section I explain four techniques I developed to aid information space design. These techniques address the comprehensibility and memorability of the structure of an information space. Comprehensibility means that structure is clearly evident. Memorability means that both the structure and how the parts fit into it are easy to remember. These design techniques make parts of information spaces more distinctive and memorable, group objects together in a space and remind users how details fit into a global structure.

Identifying places in the space with *distinctive character* is the first and most basic technique. Distinctive objects, views, and path structures make places in an information space memorable and recognizable. Thus distinctive character facilitates peripheral recognition of places in an information space, recognition that occurs without conscious efforts like closely reading the text. Peripheral recognition helps users to navigate with less conscious effort and to more easily associate memory of content to a place in the space. This lets the structure of the space act as a memory aid even after a user finishes browsing.

Character may be graphic, geometric, or structural. Examples of graphic character include colors, bitmaps, images, or animations. A picture of a chicken thumbing its nose would be an icon of interesting graphical character. Unusual geometric arrangements of common objects such as the outline of a square or a crescent moon would make that view more memorable and recognizable. Structural character refers to the structure of connections by which a user can travel between objects. A junction where paths meet is a mem-

orable place in a space because everyone must pause for a moment there to decide which way to go.¹

Second, *common character in conjunction with locality or path continuity* can bind groups of objects together into memorable structural units. To cluster together a group of objects, a designer would apply similar distinctive character to all the members. This could be as simple as painting them all the same color. A conceptual motif of graphic character could be used if it clear enough and well differentiated from the rest of the space. For instance, if items were labeled with icons depicting chickens in various states of activity, users would be likely to recognize each of the chickens as belonging to a group.² In addition, objects share a geometric character as a group if they are all members of the same geometric arrangement.

The sense of locality is then created by proximity or by the structure of interconnections. If it is easy to get to any member from another, the objects form a district. If they are connected in an ordered sequence, the objects form a path. Even if objects are structurally distant at opposite ends of a path, the whole path will still have some perceived unity.

Reminding vistas are views of the overarching structure of a site. Reminders help to connect the local experience of examining details to a more global structure. A general reminding vista shows a complete graphical representation of a structure. Maps or even text outlines are examples of general reminding vistas. A local reminding vista is an incomplete graphical hint, shown when a user is examining a detail, that suggests the global structure and where in that structure the detail fits. Well designed signage can have such an effect by pointing out the relationships between distant points in the space. Also, if a space is divided up into regions with different character, perhaps painted with a variety of textures, a glimpse of one such texture may remind users of the pattern of textures across the space.

Reminding vistas are not as effective, though, if the parts of the vista are not distinguishable from each other by character, or if they differ only in some unmemorable or

¹My demonstration in Section 2.3 does not make use of structural character because the zooming and panning interface does not include a notion of interconnecting paths on which users travel.

²No chickens were harmed in the course of this research.

random way. A common example of this would be the assignment of random colors to every member of a group. A vista of such random assignments would remind users of the global structure only in the abstract. Users would remember only that the space was divided into areas of different colors instead of how the elements fit into the global structure.

Countable character addresses the problem of how to distinguish memorably between elements of a set. It is a method of associating unique character with each of the members of a small group. Based on the number of elements in the set, a recognizable and predictable set of character traits is applied. A designer might depict a three element set by using red, white, and blue. A two element set could be depicted in black on white and white on black. In my own demonstration, I used lines drawn vertically, horizontally, up-and-right, and up-and-left to characterize a set of four members. These lines let a user who is zoomed in to see only a small portion of the space be reminded of the four divisions of the global structure and of which section is in view. Also, like a theme of common character, membership in a recognizable countable character set helps to bind together a group of objects.

It is not always easy to come up with countable sets. Also, one must be careful about sets that are only verbally memorable. For instance, it is easy to remember the mnemonic “ROY-G-BIV” which gives the colors of the rainbow in order. It is much more difficult to visualize a rainbow with all its hues in order.

Countable character helps users to monitor their progress as they browse and provides reminders of content. Some questions become easy to answer when the user can easily recognize and enumerate the members of the set: “How much have I seen?” “How much do I have left to see?” I believe that users familiar with the character set that is associated with a structure will be able to step quickly through the set in their heads to note which parts they have or have not inspected. This occasional review will help users to remember the contents of sites.

One additional way to think of countable character is to consider it to be a generic form of the memory palace technique. To remember a lengthy speech, an orator would conceive of a palace of wild and imaginative scenes where the rooms would hold parts of the speech.

The orator would then walk through the rooms of the palace in the prescribed order, recalling its contents. Countable character differs in that the “palace” is carefully designed not to rely on any private ideas of the author’s. It is not the author who steps through the “rooms” of the “palace,” but a diverse collection of users. For countable character, the “rooms” therefore must be a set of character traits that the audience will readily recognize.

2.3 Demonstration tour

Content and Organization. To demonstrate these techniques, I built a tour of some of the research projects at the MIT Artificial Intelligence Laboratory. The structure of the site is inspired by a model of A.I. Lab research used in talks by Patrick Winston.³ Information about research projects is placed in the space according to the conceptual “connections” strengthened by the projects. Projects lie between pairs of concept icons selected from four concepts: The Information Network at the bottom, People in the middle, Computers at the upper left, and the Physical World at the upper right (depicted by a rock.) For example, a vision project, helping computers to experience the physical world, would be between the icons representing Computers and Physical World. (Figure 2-1 shows the general layout of these concept icons.)

Demonstration interface and platform. The zooming and panning interface gives users an essentially infinite resolution image to see at any scale and from any vantage point. The wide range of scales gives the space depth, but unlike a fully 3D fly-through, users only look at the scene; they cannot change the direction the camera is facing. This is called a 2.5D interface.

I chose this model for my demonstration because of its simplicity, continuity, and maintenance of context. (Appendix A reviews a number of information visualization methods

³This model is reflected in a set of slides on the World Wide Web at <http://www.ai.mit.edu/director/introduction.html>

and interfaces. Zooming and panning interfaces are also discussed further in the Chapter “Future Work.”)

I implemented the demonstration space using beta releases 0.2.7 and 0.2.9 of Pad++. Pad++ is an extension to Tcl/Tk that introduces the pad widget, which is similar to a Tk canvas on which shapes and text can be drawn, but with additional facility for animated zooming and panning to view the pad from any position and scale. Pad++ comes with a MacDraw-like program for creating pads called Paddraw which I used to draw and arrange the icons and text. I also considered using a 3D walk-through interface, where users walk around a virtual landscape, seeing it from a first person perspective. I chose the zooming and panning interface because it was much simpler without much difference in flexibility.

Text in the tour. Text is an important part of the tour. To avoid clutter, to speed up rendering, and to control when passages are read, the text is animated simply by fading it in and out at appropriate times. For instance, when a user zooms in on the heading, “Artificial Muscle,” the explanatory text, “Linear actuators that undergo reversible length changes in response to chemical stimuli.” fades in at a slightly smaller font. It fades out again when the user moves on.

Example scenes and sequences. This section discusses some highlights of a demonstration tour of the MIT Artificial Intelligence Lab, showing examples of how I applied the techniques described in section 2.2.

Figure 2-1 presents the overarching structure of the tour, derived from Winston’s lecture. To be introduced to the structure, users zoom in to see animated text explanations of each of the four nodes: The Information Network at the bottom, People in the middle, Computers at the upper left, and the Physical World at the upper right (depicted by a rock.) Then users similarly introduce themselves to the connections between those nodes.

The distinctive “martini glass” shape of the overall connections map makes it useful as a general reminding vista. Seeing that shape during the tour should remind a user of the global structure of the site. This configuration also lends itself to being a local reminding




Figure 2-1: The Martini Glass: Global layout of the demonstration space.



Figure 2-2: Connecting Computers to the Physical World

vista thanks to a convenient instance of countable character. The orientations of the four connecting lines form a small, complete set: (—, /, |, and \). In the thick of the tour when the overall structure is not visible, the connecting lines remind users where they are in the global structure by their orientation. For instance, when zoomed into the projects connecting People and the Physical World, the connecting line that is visible is angled up and to the right. These lines remind users of the general structure of which they are seeing a part. They convey, to users who are aware of the general shape, where they are on the map. This helps users to attach memories of specific projects to the appropriate places on the map, making them easier to remember and find.

Zooming in on the area connecting Computers and the Physical World, (depicted in Figure 2-2) shows the local reminding vista effect of the connecting line. The horizontal, gray line indicates that this is the top portion of the map and the connection between




Figure 2-3: Cog and the Robotics Projects: The robot Cog's posture is reflected in the arrangement of the robotics projects. Graphic and geometric character reinforce one another to make this view easily recognizable.

Computers and the Physical World.

The rings and arrows in this view are graphically recognizable and interesting. Their arrangement is geometrically interesting as well. Together, these characteristics allow quick peripheral recognition of the area without examination of the text. The three ring icons, with arrows pointing in, out and both in and out, are an example of countable character. This makes the set more memorable as a group.

Finally, zooming in further to Figure 2-3, Robotics is distinguished by a picture of the humanoid robot, Cog, presiding intently over the space. The geometric arrangement of projects (each of which is a robot) reflects the outline of Cog's posture. The graphic character of the image of Cog and the geometric character of the arrangement of projects

are meant to reinforce one another to make this area distinctive.

Chapter 3

A Sense of Place

In this chapter we leave behind the sense of the space to examine how a user's sense of his or her *place* in a site is developed and maintained while browsing.

The sense of place is the way a user experiencing a site perceives his or her relationship to the structure of the site. Most simply put, it is the user's answer to the question, "Where am I?" Intuitively, the less work a user must do to maintain a sense of place, the more attention he or she will be able to devote to learning information, remembering structure, and carrying on tasks in the space. I argue that it is possible and desirable to build sites that support users' navigation in such a way that at a navigation decision point users need only a minimal, local sense of space, e.g. "I am in my office." rather than an involved chain of information, "I am in my office, which is on the eighth floor of the A.I. Lab, which is the second identical tall building on Main Street ten minutes down from Mass Ave, which is..." A local sense of place requires no memory or cognitive effort, but only simple recognition.

A good sense of the *space* as discussed in the previous chapter can have a positive impact on the sense of place. Knowing the structure of a space can encourage users to think more locally about their place. Users who understand the structural relationships between locations in the space need only to recognize their immediate location in order to find their way to another known place in the space. Users who fear getting lost want to know more information about their place in order to navigate.

Lacking a detailed sense of the space need not cause a user to feel lost, though. A

user may feel perfectly comfortable navigating an unfamiliar area if she does not expect any navigationally demanding tasks. For example, if a user knows that an unfamiliar path she is following will lead to someplace known, she will not feel uncomfortable simply to follow it. She may thus pay less attention to where she is globally with respect to the site.

In this chapter I present three interface features I developed to encourage navigation with a simple sense of place. I incorporated these features into a Web site I call the Map site. I then conducted an experiment showing that users of the Map site were not inclined to maintain a sense of place that related their place to their navigational histories. Navigation in sites without these three features did show signs of this tethered navigation.

3.1 Interface features to simplify the sense of place.

This section presents three features of interfaces to Information Spaces that are intended to engender a simple sense of place in users: a largely static interface, a continually visible place indicator, and immediate access to top-level navigation. I implemented these features in a Web site called the Map site, which is described in section 3.4.3. I ran an experiment (discussed throughout this chapter) which showed that the Map site did encourage users to disregard the route they came by when navigating. In other sites, subjects showed signs of awareness of this history.

The first interface feature is a *largely static graphical interface*. This is to say that many components of the interface do not change as a user browses. People browsing in normal hypertext feel as though they are immersed in the space, personally moving through it. A static interface is less of an immersive inhabitable world and more a display that can be viewed and manipulated. Users will feel less like they are immersively “walking through” and more like they are “reaching into.” Without being so personally invested in the space, I believe that users will be less apt to fear getting lost and thus less likely to concern themselves with how to get back. It will be less of a shock to simply pick up and start again. Thus users will be able to give more attention to scanning, finding, and learning

information. To experiment with a static graphical interface, I divided the Map site into Netscape frames where browsing was relegated to one small frame.

The second interface feature to encourage users to attend less to their global sense of place is a *continually visible place indicator*. Users should trust that they are free to forget their place in the global structure. Besides visual recognition as discussed in Chapter 2, a constant indicator in the interface may help users to “lose” themselves in the site. If not always visible at a glance, this feature should at least be accessible without having to leave and come back to the current place in the space, perhaps by using a pop-up window. I used a map with a “You-are-here” indicator to test this feature.

This map additionally shows the general structure of the site. Having that information always present on the display could have one of two opposite effects on the user’s sense of place. It may help users to learn a sense of the space, which, as I have argued, decreases the need for a complicated sense of place. On the other hand, off-loading this global, structural information to the interface may encourage users not to bother to comprehend and remember the structure of the space. My experiment is not precise enough to determine which is the case. Though it would be interesting to determine which is the case, this ambiguity does not diminish the experimental result that these three features do encourage a more highly minimized sense of place.

The last feature is *immediate access to top-level navigation*. I believe that much of the reason that users keep track of their path behind them is the need to feel that they can safely get back to a home node, a central page from which they can easily get to many places they may want to visit. Immediate access to top-level navigation infuses the role of the home-node into every page of the site by making it inherent to the interface. I effected immediate access to top-level navigation in the test site by making the continually visible map active such that it could be used to get to any location by simply clicking. No page in the site was more central to the top-level hierarchy (and therefore more home-like) than any other.

3.2 Behavioral manifestations of the sense of place

I hypothesized that the three features described in the previous section would produce a site where users would browse with a more local sense of space, less conscious of their way back to a home node. To test the effects of these interface features experimentally, I had to quantify the degree to which users were keeping track of their path behind them. Of course I could not directly quantify a subject's mental state, but I could measure a behavioral manifestation of the subject's desire to keep track of the path behind them. I inferred that those users who use the "Back" button to return to the top navigation level of a web site are more intent to remember their route than other users, who make less use of the "Back" button. Other users would use interface tools that would bring them immediately to the top level. I took this to imply that these users were less attached to their entry route, and more comfortable striking out on a new but more direct path home. By comparing the behaviors of groups of subjects browsing in different sites, I sought to get an impression of the effect each site has on its users' senses of place.

What does the "Back" button indicate about users? One obvious conclusion is that they are aware of where they've been. Navigation relying on the "Back" button requires users to keep track of the steps they made to get to where they are as far back as they might want to back out. "Back" button users are also likely to be keeping track of distances back to important points. Otherwise they couldn't make quick decisions about whether it would be easier to back out or to try a more complicated interface tool like typing a URL, or a less ubiquitous and less familiar tool like a navigation button particular to the given hypertext.

Navigational reliance on the "Back" button can cause people to use it even when it is not efficient. If a user finds the "Back" button useful for navigation, then he or she will feel pressure to *maintain* the browser's stack of locations. Users will actually develop habits of pruning their browser stack to get rid of tangential explorations and pages they don't need to see again. This pruning takes place by moving back through the site even if that is not an

efficient or sensible way to continue.¹ Thus the history stack—which seems like it should be a helpful memory aid—can actually be a burden on both memory and navigation. A user who wants to stay within reach of nodes behind her must not only keep track of the useful places she has been, but she must also keep track of and regularly purge all of the useless places she has visited.

3.3 Basic format of the experiment

To test subjects, I first asked them a question about the context of the site, thus prodding them to browse in a sub-hierarchy. After they had answered the first question, I asked a second question to bring them out to the top level. Subjects had the option of clicking a link to take them directly to the top level, leaving their recent navigation history on the browser stack, or of using the “Back” button to back out to where they had last navigated in the top level, cleaning up the history stack. “Back” button usage at these times indicates a sense of place which connects those subjects to home. I repeated the question for a total of thirteen questions. Finally, I compared the groups of subjects that used each site to see what effect each site had on its users in this regard.

I demonstrated experimentally that subjects often use the “Back” button to get home in traditional web sites even when other options were available. On the other hand, I created the Map site (section 3.4.3), using the features of Section 3.1 where not one user clicked the “Back” button to get home. This meant that the site supported those users to navigate without having a detailed sense of place tying them back to home.

¹ An interesting project would be the modification of a browser so that a user could edit the history stack without moving.

3.4 Sites built for the experiment

In the creation of the web sites for this experiment, my goal was to create browsing experiences with clearly, qualitatively differentiated senses of place. I built three Web Sites: the Standard site, the Map site, and the Disjoint site.² I built the Standard site simply to show contrasts with the other two sites. The Map site was designed to incorporate all of the features discussed in Section 3.1: a largely static interface, a continually visible place indicator, and immediate top-level navigation. It was my intention that users of the Map site navigate with a local sense of place rather than a continual awareness of where they came from. I additionally created the Disjoint site to see if I could show a further contrast and increase users' dependencies on their histories by implementing the inverse of two of the features. I did not find a significant difference between the Disjoint site and the Standard site. However users' behavior in the Map site was clearly differentiated from the other two.

3.4.1 The Standard site

The Standard site (pictured in figure 3-1) was intended as a control in the experiment to suggest the current state of the World Wide Web for comparison. This site was a slight modification of the Artificial Intelligence Lab's actual research page.³ The top-level page of the Standard site listed a subset of the projects from the Lab's home page. Each project had a brief description copied directly from the Lab's page. The projects were grouped into categories in the same way as on the A.I. Lab homepage. The size of the groups varied widely from just one to thirteen projects.

Each project's name was hyperlinked to that project's Web site. I did not alter these project sites except to insert a link at the top of each page to jump "Back to the top level." (An internal page modified with a "Back to top level" button is shown in Figure 3-2.)

The graphic display of the Standard site was not static at all. Besides the constant "Back

²These sites are not publicly available for perusal on the World Wide Web because in their implementations I made use of the fact that the server and client were running on the same workstation

³<http://www.ai.mit.edu/projects.html>




Figure 3-1: The top-level page of the Standard site contains links to sites about various projects with brief descriptions.




Figure 3-2: In the Standard and Disjoint sites, the “Back to top level” button is the alternative to using the “Back” button to get to the top level.




Figure 3-3: The Disjoint site is a deeper hierarchy than the Standard site. It contains no visual or conceptual clues to help a user find a project from the top level

to the top level” button, the whole screen changed when a user follows a link or backed up. A view of the top-level structure and access to top-level navigation were both indirect. To view or use the top-level page, users had to leave the page they were on. There were no cues in the site to show users where in the top-level structure, that is, in what group of projects the project site they were browsing was found.

3.4.2 The Disjoint Site

The Disjoint site (pictured in Figure 3-3 and 3-4) was intended to *increase* users’ dependence on their browsing history still more than the Standard site. This site contained the same projects as the Standard site, but the projects were divided randomly into three pages




Figure 3-4: The second layer of hierarchy in the Disjoint site has three pages like the page depicted here listing randomly grouped projects.

each containing four or five projects. The top-level page opaquely linked to “Page 1,” “Page 2,” and “Page 3.” Like the Standard site, the Disjoint site also included “Back to top level” buttons on all of its pages including the three pages of projects just under the top-level page.

In its immersiveness, the Disjoint site was similar to the Standard site. The user’s view changed completely with each step rather than having static areas. The Disjoint site provided even less accessible visible indication of the user’s location than the Standard site. Users could never view the whole top-level structure at one time. Instead, they had to synthesize the three pages using memory. Of course, the opacity of the top-level page, the random groupings and the lack of associations of character as to structure discussed in Chapter 2 meant that users of the Disjoint site were not aided in any way to memorize the locations of the projects. Finally, navigation through the top level was less immediate than in the Standard site because of the two level top-level hierarchy.

3.4.3 The Map site

The Map site is my effort to demonstrate the features discussed in Section 3.1: a static graphical interface, continually available location information, and continually available top-level navigation. The static graphical interface was achieved by dividing the browser window into frames, only one of which would change during normal browsing. A graphical map of the top level of the site with a “You-are-here” push-pin in a static frame provided continually available location information. The map was also active; clicking on a project shifted the browser to that project. Thus the map also provided continually available top-level navigation.

The Map site makes use of HTML frames to divide the interface into three sections. (The whole interface is shown in Figure 3-5.) The “Overview” frame is at the lower left, the “Browsing” frame is at the lower right, and the “Top-level Navigation” frame is the largest frame at the top.




Figure 3-5: The Map site is designed to minimize users' dependence on their browsing history and to encourage a local sense of place. The frame structure creates a *largely static interface*. The map in the upper right provides a *visible place indicator* and *immediate top-level navigation*.

Overview frame. An overview of the current project, derived from the project's home page, is displayed in the lower left frame. Clicking on links in the Overview frame puts the linked page in the lower right Browsing frame; the Overview frame does not change unless the user visits a new project.

Browsing frame. Users browse Web pages taken directly from A.I. Lab projects' web sites in the Browsing frame, on the lower left. When a visitor arrives at a new project this frame is initialized with a menu derived from the project's home page. This menu contains links to various project pages such as lists of projects, people, or papers. A link labeled "Back to top level" at the top of each project's Overview frame restores the Browsing frame to this initial menu. This frame was intentionally made small compared to the rest of the interface to test the effects of a largely static interface. To prevent subjects from having to scroll quite so often browsing in such a small window, I stretched the Netscape window to almost the full span of the 17 inch monitor. Though a few subjects mentioned wasted screen space in the top navigation frame as a design flaw of the Map site, no one directed comments at the browsing window itself. Therefore I believe the window, though small, was sufficient for browsing. These unsolicited comments were offered during the debriefing after the experiment.

Top-level Navigation frame. The Top-level Navigation frame contains the name of the current project printed in blue, along with a brief description. There is also a message here saying which pair of concepts is "connected" by the project. Here connections have the same meaning as in the demonstration tour in Chapter 2. In addition there is space reserved for relevant technology and application buzzwords like "vision" or "mars exploration." In hindsight these buzzwords were a poor usage of screen space. Many projects had no interesting buzzwords to include, and there was not enough space to fit useful explanations of each term. The buzzwords were also out of place in a frame devoted primarily to top-level navigation.

The Map. The most noticeable feature of the Top-level Navigation frame is a graphic map of the projects which is always visible on screen. The map is laid out in the same “martini glass” pattern as the zooming and panning demonstration tour discussed in Chapter 2. Projects are grouped into four clusters making connections between People, the Information Network, Computers, and the Physical World. Each project is represented by a blue dot on the map. The current project’s dot is pierced by a green “You-are-here” push-pin. If a user drags the mouse over any blue dot or over any of the four concept icons, the Netscape status bar at the bottom of the screen indicates the name of the project or concept. A click on any of these dots transports the user to that project, updating all three.

A pink path winds through all of the four major concepts and the projects “connecting” them. To the lower left of the map is a pink area with buttons for jumping to the next and previous projects along the path, as well as a small, additional map to show what directions the “Next” and “Prev” buttons will move the user on the Main map.

Although its most salient feature is a map, the Map site is spatial in only a limited sense. Space is used on the map for visual clustering. Projects have locations on the map⁴, and they are grouped by their placement between the same concepts. Also, users may understand their positions spatially with respect to the map. For example, one might think, “I’m up at the top of the map near Cog and the other robots.” However, motion around the map in the Map site is not immersively spatial like motion in a city or building. Aside from the small Browsing frame users see themselves from an external perspective on the map.

3.5 Procedure

This section documents the sequence of events of the experiment. Subjects were first introduced to the three windows used in the experiment: The Questionnaire, the Netscape

⁴It has been argued that giving objects a location makes them easier to remember. Schulman [Sch83] presents a positive argument which is contradicted by Jones and Dunais [JD86] [Yat66] gives a detailed historical perspective on memory techniques, many of which employ spatialization. Thus the localization of projects may have aided subjects in remembering the top-level structure.




Figure 3-6: A typical screen configuration during the experiment, including the Questionnaire, Netscape browser, and “Back” button.

browser, and the “Back” button in its own window. (A typical screen including these windows is shown in Figure 3-6.) I first introduced subjects to the interface. For the Map site this included pointing out the four main concepts and their icons, Information, People, Computers, and the Physical World. In all cases I pointed out the “Back” button (a Tcl/Tk button in its own window) and the Questionnaire (another Tk window.) I told users they could click on the “Back” button I provided or use the alt-left key combination, which I intercepted with the window manager for logging. I could not find a way to intercept or prevent use of Netscape’s right-button menu, another way to back up, so I simply asked subjects not to use that feature. I ran the experiment using Netscape 3.01 with all optional tools hidden: the menu bar, location, toolbar, and directory buttons.

After my introduction the Questionnaire took over, guiding subjects through the rest of the experiment which was divided into three parts. First, subjects were asked how often




Figure 3-7: Instructions given to the subjects as they begin to familiarize themselves with the site.

they used the World Wide Web and what they used it for. This was the only personal information collected about the subjects. Second, subjects were asked to take about 15 minutes to familiarize themselves with the site so that they would be able to find answers to questions when asked later. Figure 3-7 shows the text introduction in its final form⁵ as displayed by the Questionnaire for this phase of the experiment. Most subjects took 20 to 25 minutes browsing because they got caught up in reading about projects they found interesting. For those that took the longest I hinted that they should watch the clock and move on, but I did not enforce a strict time limit.

In the final phase, the Questionnaire asked the subjects a series of 13 questions. After each question was asked, the subject had to get to a new sub-hierarchy via the top level in order to find an answer. The questions had keywords planted in them to make it easier for subjects to find the answers. For instance, to find the answer to “Which Leg Lab robot was inspired by the kangaroo?” a subject must go to the top level, click on “Leg Laboratory,” click on “Our Robots,” and then finally read down the list of descriptions of robots. Many answers were more difficult to find than I had anticipated. Subjects often needed to look in a few other sub-hierarchies before they found the right project. Finally, after the subjects were finished finding answers to the questions, I debriefed them to explain the methods and purpose of the experiment.

I found that I had to make some changes to my introduction in the middle of the experiment. After I noticed that some subjects had gone off-site by clicking on a link to the Artificial Intelligence Lab’s home page, I began warning new subjects that the A.I. Lab homepage should be considered off-site. I added that warning into both the Questionnaire’s written introduction and my own verbal introduction. Later, one subject from the Standard site told me after his debriefing that he had thought my instruction about the A.I. Lab homepage applied to the “Back to top level” button as well. He used only the “Back” button during the experiment. Realizing that my new instruction was being misunderstood, I

⁵I had to make changes to this introduction after some subjects had already participated to keep subjects on track. These changes are explained below.

began pointing out in my verbal introduction the difference between the “Back to top level” button and the ubiquitous links to the A.I. Lab homepage. (His data set is not included in the results.) Three other subjects used only the back button, it is possible that they may have had the same misunderstanding. I did include the data for these questionable subjects in my analysis.

I made one other change when I discovered that some subjects were trying to use the START natural language system to answer questions. The START page, which could be found on all three sites, had a demo capable of answering questions about the A.I. Lab posed in English. When I discovered this I additionally asked subjects not to use the START system to answer the experiment questions. I put this warning into both the written and verbal introductions.

Subjects. There were 32 subjects. Eleven browsed the Map site, eleven browsed the Standard site, and ten browsed the Disjoint site. Three subjects from the Standard site and two from the Disjoint site were disqualified from the experiment. One misunderstood the instructions. The other four left the experimental web site during the experiment. Reasons for disqualifications are discussed in more detail in Section 3.6.

Subjects were recruited by posters on MIT’s campus and through word of mouth. They were compensated ten dollars per hour for about forty-five to ninety minutes of participation which took place in my office.

System. I ran the experiment on a SPARC5 running SunOS 4.1.4. All the subject’s actions were logged by the CERN proxy server (version 3.0A) running on the local workstation. The Webfilter patch to the CERN httpd (beta version 0.5) gave me the ability to insert “Back to top level” buttons by filtering the HTML with my own Perl scripts. When the Questionnaire asked a question or recorded an answer, that text was merged into the web server’s access log. The window manager twm intercepted the alt-left key combination. When a subject clicked my “Back” button or pressed alt-left on the keyboard, the string, “Back!” was merged into the log before a system call of `netscape -remote 'back();'`

prompted the browser to back up.

3.6 Data collection and analysis

Relevant data was recorded after a new question was asked. At that point subjects had to make a transition to the top level from the sub-hierarchy they were browsing. Then they would either back out with the “Back” button or use a top level navigation interface. In the Standard and Disjoint sites the alternative to the “Back” button was a “Back to top level” button. In the Map site, subjects could immediately begin navigating in the top level by clicking on a project on the map or by using the “Next” or “Prev” buttons to move along the trail connecting the projects.⁶ Because it supported immediate top-level navigation, backing out in the Map site would not have made available any other top level information or navigational tools. It only would have served to clear the browser’s stack. For this experiment, it was irrelevant (and I did not record) whether subjects used the “Next” and “Prev” buttons or clicked directly on the map. Either way, they were able to abruptly leave behind their browsing history to navigate in the top level.

To compare the sites I ranked all the subjects according to how much they used the “Back” button when going back to the top level for a new question. For each user I gave each question a score indicating whether the Back button was used. If the “Back” button was used exclusively for a question the score was one; if not at all the score was zero. I averaged the scores for each subject’s questions to get a score for that subject.

In some cases users backed up part-way and then jumped back to the top with the “Back to top level” button. In those cases I counted the number of Back-clicks that would have been necessary to get back to the top level and scored those questions the ratio of the number times the “Back” button was clicked to the number necessary. This case occurred

⁶At risk of confusion to readers, the “Back to top level” button in the overview frame of the Map site is not a top-level navigational tool like the buttons with the same name in the Standard and Disjoint sites. The Map site’s “Back to the top level” button only takes the user to the top level of that project, not the top level of the whole site.

Table 3.1: Table showing usage of the “Back” button in each site to get back to the top level. A subject’s score is the fraction of the number of trips back to the top level where he or she used the “Back” button to get there. In the Map site the results were absolute: nobody used the “Back” button. Results from the other two sites were less clear, but they definitely showed “Back” button usage.

Site	Number of subjects whose “Back” button usage scored $\leq$										
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
Standard site	-	-	-	1	-	2	1	-	2	1	2
Disjoint site	-	1	3	1	-	-	-	2	-	-	1
Map site	10	-	-	-	-	-	-	-	-	-	-

only in the Disjoint site.

Some questions that were asked did not bring the subject from a sub-hierarchy to the top level. This simply decreased the number of data points being weighed into a subject’s score. I ignored questions where the subject did not go to the top level because the previous question left them in the same sub-hierarchy. I also ignored questions where the subject was already at the top level when he or she finished the previous question. Subjects would occasionally go back to the top level in anticipation of the next question before they clicked “Go on” to see what the question was. Also many questions in the Map site could be answered from the overview without the subject leaving the top level. Finally, I ignored subjects’ actions in response to the first question. I wanted all subjects to be performing the same question answering task leading up to the data recording, as opposed to the free familiarization task that preceded the first question.

I also had to throw out the data from some subjects. Four subjects left the site I built and were instead using the main A.I. Lab site. Those subjects were no longer browsing the site about which I had intended to learn. A fifth subject told me after his debriefing that he thought he was not allowed to use the “Back to top level” button.

Table 3.1 shows how much different subjects made use of the “Back” button in each site. A higher score shows more dependence on the “Back” button. The skewed data recorded for the Map site made it clear that the site was used differently than the Standard and Disjoint sites. No subject in the Map site ever used the “Back” button to clear the

history stack before moving in the top level.⁷ Contrarily, no subject in either the Standard or Disjoint sites ever went through the whole experiment without using the “Back” button to get to the top level.⁸

3.7 Discussion

Explanations of the data. The Map site made use of the three features of Section 3.1 to achieve this success in freeing subjects from their navigational history. Because I tested only the conjunctive hypothesis, there are three interpretations of this success. It may be that the three features, the largely static interface, the continually available place indicator, and the immediate top-level navigation all supported users to some degree in simplifying their sense of place. It may be that only one or two of the features had a positive effect. It may be that the variety of features addresses the various navigational styles of a diverse population of subjects. Although a finer-grained experiment would be required to separate these interpretations, I count the experiment on the whole to be a success. I have shown that the three features, static interface, ever-present structure and location reminders, can be used to predictably encourage users to navigate with a simpler sense of place.

There are other reasons why users might have been less inclined to use the “Back” button in the Map Site. One reason is that the behavior of the “Back” button in frames-based sites like the Map site has changed since Netscape first released a frames-capable version of its popular browser. When using Netscape version 3.0 the “Back” button takes one to the previous configuration of the frameset on screen. Netscape 2.x, however, took

⁷Though it was never used to get to the top level, the “Back” button was, on rare occasions, used in local browsing in the Map site. As a quick metric, I scanned the logs for the total number of Back clicks in each site using `grep`. This scan included logs for subjects and questions that were rejected in the actual analysis of the experiment, and probably includes extra “Back” clicks that may have occurred while I was preparing the experiment for each subject. In the Disjoint site, there were 543 recorded “Back” clicks for 10 subjects. In the Standard site, 540 such clicks were recorded for 11 subjects. In the Map site for 11 subjects, there were only 24 “Back” clicks.

⁸One subject in the Disjoint site came extremely close. That subject never used the “Back” Button except for one question, even then clicking it only once before skipping the four remaining steps with the “Back to top level” button.

the entire frames-based site to be an atom on the history stack. Even if a user followed many links within a frames-based page, the “Back” button would take her back to the page she left to get to that frame site, throwing away the state of where she had browsed in that site. Thus a large class of users who first used frames with the older and more frustrating “Back” button behavior may be afraid to use that feature when they see frames. Worse yet, they may be unaware of the difference between the versions and simply find the behavior of the “Back” button within frames to be completely unpredictable.

One subject in the Map site said that he did not use the “Back” button because he did not treat the space like a World Wide Web site. The space felt to him like an entirely different kind of program because it lacked the usual web browser toolbars and buttons, and because the interface with its large static area was unusual. The lack of toolbars was common across all three sites, but the other sites did look more like traditional web sites that would be familiar to Web users.

Some other factors may have influenced “Back” button usage but should have had the same effect on all three sites, and therefore these factors should not have influenced my comparison. Some subjects were aware that Netscape was optimized to load pages faster with the “Back” button than through normal hyper-links. Also, the “Back” button was the only interface tool to have a keyboard accelerator. I could have avoided this factor by asking subjects to use only the actual “Back” button and not the keyboard accelerator. Another factor is that the “Back” button is effectively part of the interface of *every* World Wide Web site. Most if not all subjects presumably came to the experiment knowing how to use it. This factor was somewhat abated by the fact that the “Back” button was not in its usual place and looked different than the subjects would have expected. Everyone had to learn the “Back to the top level,” though. Finally, the ease of being able to click on a single interface tool repeatedly may have encouraged subjects to use the “Back” button even if they were deep enough in the site that the number of clicks would have been a factor. Any of these effects could have been felt similarly across all three sites, however, and should not have affected my results.

Criticisms of the experiment. One criticism of the execution of the experiment is that I changed the introduction to the experiment a couple of times over the course of the experiment to repair some misconceptions and ambiguities that were causing subjects to do the wrong thing. More beta-testing of the experiment may have helped this, although I did not discover the problem until perhaps half of the subjects had already participated. Also, the experiment would have been more stable and precise if I had eliminated as much as possible of the verbal introduction, leaving only the written introduction provided by the Questionnaire.

Also, many of the questions in the Map site could be answered from the initial detailed overview without clicking in any further. Because these answers were so shallow in the site, I had substantially less data about users in the Map site than in the others. Scores for Map site subjects were generated from between six and as few as three data points showing a choice made of how to get to the top level.

I pointed out early on in this chapter that techniques for clarifying the sense of the space could have additional beneficial effects on the sense of place. In my design of the Map site I did not focus on these techniques, but instead relied naively on the work I had done in Chapter 2 by transferring the graphical layout to the new interface. Many of the characteristics that were effective on the zooming and panning interface did not make sense in a static map. Certainly, local reminding vistas could not exist in a map that was always completely visible and static. Also, without seeing the parts separately it is questionable that users would recognize the countable character set of line orientations inherent to the martini glass design. Indeed, the general thrust of the techniques in Chapter 2 was to give users peripheral reminders of the structure unobtrusively throughout their exploration of the site. The static map is not meant to be looked at often, but rather trusted to be there when it is needed. Any such reminders inherent in its graphic design will have few opportunities to remind.

Though the experiment is open to criticism, I believe it has a great deal of significance because it measured a specific user behavior. It is important in researching information

interfaces to identify specifically what the hypothesized effects of a feature are, and if at all possible to test them thoroughly in controlled experiments. Intuition, philosophy, and radical creativity are all useful and necessary for inspiring new ideas for information presentation, but they are not sufficient for validation of those ideas. Asking a sample of users whether they like an interface is certainly important to know, but it does not help designers to understand the specific effects of an interface such that they can determine for what tasks and situations it is helpful.

Chapter 4

Conclusions

The sense of the space and the sense of place are useful concepts to help focus the design of information spaces. They help to address the limited perspective of users who must incrementally synthesize knowledge about a space. For example, sites may be created with sensible organizations that a user may never understand by browsing the site. Useful interface tools may be provided that users may avoid for fear of getting lost. The techniques presented in this thesis can be applied to make the structure of information spaces more comprehensible and memorable, and to give users confidence to navigate freely without paying as much attention to details of navigation.

Distinctive character, common character, reminding vistas, and countable character are methods of “painting” the structure of a site to make structural relationships stand out in a coherent way. They aid the recognition, comprehension, and memorization of structural forms.

Static graphical interfaces, continually available place indicators, and immediate top-level navigation are interface features that can help users to navigate without occupying their minds with their browsing history. A user’s attention is freed up for the activities and the information for which he or she came to the information space.

Finally, it is important to not only learn what is interesting and appealing to users, but to see how a site affects a user’s behaviors when performing tasks in the site. Much can be learned about how users perceive a site and their relationship to it by watching their patterns of behavior while using it. Behaviors such as “Back” button usage can be used as

heuristics to hint at the way users perceive an information space.

Chapter 5

Future Work

As I explained in the Discussion section of Chapter 3, the experiment in that chapter tested a conjunctive hypothesis. Its success was useful as a preliminary result, and it leads encouragingly to finer grained hypotheses that should be tested with more precision. It would be enlightening to build and test additional sites individually isolating the qualities of the Map site: the static interface, the continual location reminder, and the immediate top-level navigational tool.

The techniques in Chapter 2 also suggest further experimentation. The reminding effects of countable character for example suggest memory priming experiments. An experiment could test if subjects are quicker to remember content from parts of a site that are not in view but are in the same countable set as a detail that is in view.

Another step toward making these ideas more applicable is to apply higher standards of design along with the techniques of Chapters 2 and 3. The Pad tour, as well as the Map site, could be refined not only to demonstrate the principles of this thesis, but also to be well-designed, usable sites in their own right. Raising the standards of design would put these principles to more of a test.

It could be fruitful with little technological effort to extend the sophistication of a hypertext browser's history mechanism. The current notion of a stack is limiting. A browser could allow users to edit the history stack to remove items they do not need without moving

to a new page. A browser could contain multiple stacks with different “roots.”¹

With a bit more sophistication to the interface, a browser might provide a softer, less permanent bookmarking mechanism. Bookmarks that are not marked as stable would fade from prominence in the interface as new bookmarks arrived. This would retain the user’s control which is inherent to bookmarks, but would allow users to be more liberal in their selections. Users could mark any pages they might want to return to without fear of losing important pages in the noise. (One idea which would require more sophistication is a visualization of the recent browsing history with selective emphasis to highlight important pages.) All of these browser modifications are risky because they add complexity to the most basic site. The “Back” button is a trivial interface tool. These ideas trade complexity in common browsing interactions for flexibility in being able to find and retrieve old pages.

The zooming and panning interface is interesting for further research because while the interface to creating spaces can be almost as simple as drawing, the continuity of the animation helps users to keep track of where things are in context. Zooming and panning interfaces need a better basic toolkit of navigational methods. It takes too much effort for users to move from one view of the space to another. Users do not benefit by being able to move in continuous gradations to views of a static space. It is also difficult in a zooming and panning system to get a sense of what the current scale is. Techniques like those in Chapter 2 could be adapted to aid users in recognizing the scale of their view.

Walk-throughs also provide continuous animation, and they alleviate the confusion (along with some flexibility) of the continuously varying scales of the zooming and panning interface. It is tempting though to think that they are a superior interface simply because of their realism and their similarity to everyday navigation in physical environments. While they may engage evolved, well-practiced navigational abilities of human beings, I am not aware of any studies demonstrating that they promote better memory, better understand of the space’s structure, or easier or faster navigation. It may be that users simply reconstruct a more structured sense of the space from their navigation which could be presented in

¹Some Netscape users simulate this by opening multiple browser windows simultaneously.

a much more direct way. Perhaps immersed navigation is a challenge we do not need in information space. Walk throughs should be compared with trivial interfaces like gigantic, static maps.

Appendix A

Review of Information Presentation Techniques

This appendix describes metaphors and tools for organization, presentation, and traversal of information spaces. These ideas described here do not directly address the design of information spaces. I felt it useful, though, to understand and evaluate them with respect to their effects on the comprehensibility and memorability of the information spaces.

A.1 Organization of information

Hierarchical directories and algorithmic layout are two ways of placing collections of objects in an information space.

Hierarchical directories. Hierarchies like Yahoo¹ break up a collection of information into homogeneous substructures. If the hierarchy is clear and predictable they are useful for retrieving documents. The stepwise motion from level to level can be tedious, though. The homogeneity of these directories is their biggest drawback for the memorability of structure. Without reading the text, each location in the hierarchy looks like all the others. There is no sense of location and no sense of direction of motion.

¹<http://www.yahoo.com>

Algorithmic methods for laying out information in space. Multidimensional Scaling (MDS) is a well-established algorithm from the 1950's [Tor52] that spreads objects over a Euclidean space to cluster them by similarity. Trajectory Mapping (TM) is a much newer technique based on interpolations and extrapolations of the relationships between pairs of objects in the information space. Based on this data, TM points out the sensible paths through a space organized by MDS. [RK94] Thus TM makes up in part for the reduction in dimensionality of MDS by adding the path information. These algorithms are a useful research tool for testing a population to see their preconceived notions about the information. This information is useful for organizing information in ways that users will find predictable. These algorithms are also useful for organizing homogeneous spaces such as collections of photographs or sound clips.

A.2 Graphical presentation of an information space

Cone Trees, graphical compression, techniques and the simultaneous display of multiple image layers can increase the amount of information that can be displayed comprehensibly on a screen of fixed size.

Cone Trees. Cone Trees are a 3D visualization technique for hierarchies. Users can rotate the tree to see other parts of the hierarchy or prune the hierarchy or look only at sub-hierarchies. [Rob91] Cone Trees conserve screen space when examining hierarchies, and their smooth animation helps to prevent users from getting confused when moving from one part of the hierarchy to another. However, Cone Trees are like hierarchical directories in that they demand a homogeneous structure. By not encouraging differentiation, they do not help users remember details and where they fit.

Two-dimensional graphical compression. Distortion oriented techniques [LA94]² vary image compression over a large two-dimensional information space according to a mathematical function (for instance, a Gaussian or a step function) so that a focus area shown clearly and the rest is squeezed into the remaining portion of the display. These techniques allow a user to scan around a large document while maintaining context. A related project, the Information Mural, uses graphical techniques like shading and anti-aliasing to spatially compress a large image. [JS95]

Transparency and Blur The rendering of transparent graphical layers with selective focus on individual layers allows a lot of information that is not immediately relevant (but could become important) to be displayed in an undistracting way. It has been found experimentally to be an effective way to present a lot of information accessibly without making the most important parts illegible [CS91] [Lie94].

A.3 Traversal of complex information spaces

Other techniques such as zooming and panning interfaces, or 3D fly-throughs or walk-throughs, add a metaphor of inhabitation and motion to the graphical presentation of information.

2.5-Dimensional Zooming and Panning Interfaces. Zooming and panning systems, typified by Pad++ [BH95], treat the world as a flat surface that gains depth when some details are much smaller than others. By zooming in on a large object or area, one finds the smaller places subsumed beneath it. Pans across the space to destinations at the same scale are accompanied by a zoom out and back in to maintain context. Zoom and Pan interfaces remain simple while maintaining continuity and context in the space.

²There is so much work in distortion oriented presentations that a review paper [LA94] was recently published to coherently identify the distinctions between various approaches.

Pad++ has inspired some novel document layouts such as a sequential ring of pages with a large title page in the middle, or a page of text that is squeezed into an inverted triangle such that lines lower down toward the point shrink exponentially compared to those above. [MPBH95] These arrangements are successful because the documents' heads are clearly visible from afar, and the text is comfortably readable within the space. The ring has the added benefit that as a user pans to the next page, the direction of motion indicates how far he is in the document. Thus in this scheme, the interface takes on the effort of keeping the users sense of place.

Galaxy of News: 3D fly-through. The Galaxy of News lets users fly through an abstract 3-space of rendered text articles. The space is organized as a pyramid with the top oriented towards the user. Documents are spread out algorithmically in the plane of the base of the pyramid with a field of large news headings above them forming a soft hierarchy. [Ren94] Again, smooth motion helps users to keep their bearings. However, because flight through the Galaxy is so unconstrained, it is easy to fall off the map and get lost, having to look around helplessly in space for a visual cue.

Urban Cyberscapes: 3D walk-through. Urban Cyberscapes is a walk-through city-space of documents with landmarks, neighborhoods grouping related information, and a "monorail" system to quickly bring users to important views. [Ull95] By providing a ground as well as the monorail, Urban Cyberscapes gives users fewer opportunities to get lost than the unconstrained Galaxy.

A.4 Search engines

One other mode of information presentation deserves mention here. Search engines like Lycos³ or Alta-Vista⁴ hide an information space behind an agent that returns information

³<http://www.lycos.com>

⁴<http://www.alta-vista.com>

objects in response to queries. Extreme dependence on them would suggest that information space design is irrelevant. Search engines are a popular and effective way to find information on the Web, but they provide no structure, leaving users to fend for themselves in understanding and remembering relationships among the contents. Also, they provide no hints about their scope aside from perhaps a boast of size.

Bibliography

- [App69] Donald Appleyard. Why buildings are known: A predictive tool for architects and planners. *Environment and Behavior*, pages 131–156, December 1969.
- [BH95] B.B. Bederson and J.D. Hollan. Pad++: A zoomable graphical interface system. In *Proceedings of ACM Human Factors in Computing Systems Conference Companion (CHI '95)*, pages 23–24, 1995.
- [CC81] William G. Chase and Michelene T.H. Chi. *Cognitive skill: Implications for spatial skill in large-scale environment*, chapter 6. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1981.
- [Cha83] William G Chase. *Spatial Representations of Taxi Drivers*, pages 391–405. In [RS83], 1983.
- [CS91] Grace Colby and Laura Scholl. Transparency and blur as selective cues for complex visual information. In *SPIE Conference*, March 1991.
- [JD86] W.P. Jones and S.T. Dumais. The spatial metaphor for user interfaces: Experimental tests of reference by location versus name. In *ACM Transactions of Office Information Systems*, volume 4, pages 42–63, 1986.
- [JS95] Dean F. Jerding and John T. Stasko. Using information murals in visualization applications. In *UIST*, pages 73–74, November 14-17 1995.
- [Kui78] B. Kuipers. Modelling spatial knowledge. *Cognitive Science*, (2):129–153, 1978.

- [Kui82] B. Kuipers. The “map in the head” metaphor. *Environment and Behavior*, 14(2):202–220, March 1982.
- [LA94] Y.K. Leung and M.D. Apperley. A review and taxonomy of distortion-oriented presentation techniques. In *ACM Transactions on Computer-Human Interaction*, volume 1, pages 126–160, June 1994.
- [Lie94] Henry Lieberman. Powers of ten thousand: Navigating in large information spaces. In *Conference on User Interface Software Technology*, Marina del Rey, California, November 1994.
- [Lyn60] Kevin Lynch. *The Image Of The City*. Mit Press, Cambridge, MA, 1960.
- [McA91] David McAllester. Observations on cognitive judgments. AI memo 1340, MIT, Cambridge, MA, December 1991.
- [MPBH95] J. Meyer, K. Perlin, B. B. Bederson, and J.D. Hollan. Two document visualization techniques for zoomable interfaces. In *Proceedings of ACM Human Factors in Computing Systems Conference Companion*. CHI, 1995. submitted.
- [Ren94] Earl Rennison. Galaxy of news: An approach to visualizing expansive news landscapes. In *ACM Multimedia*, 1994. Demonstration Proposal.
- [RK94] Whitman Richards and Jan J. Koenderink. Trajectory mapping (‘TM’): A new non-metric scaling technique. A.I. Memo 1468, MIT Artificial Intelligence Laboratory, 1994.
- [Rob91] George C. Robertson. Cone trees: Animated 3D visualizations of hierarchical information. In *ACM CHI ’91 Conference on Human Factors in Computing Systems*, 1991.
- [RS83] D. Rodgers and J. Sloboda. *The acquisition of symbolic skills*. Plenum Press, 1983.

- [Sch83] Arthur I. Schulman. *Maps and Memorability*, pages 359–367. In [RS83], 1983.
- [Tol48] E.C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55(189), 1948.
- [Tor52] W.S. Torgerson. Multidimensional scaling I. theory and method. *Psychometrika*, 1952.
- [Tve93] Barbara Tversky. Cognitive maps, cognitive collages, and spatial mental models. In *Proceedings: Spatial Information Theory*, number 1, pages 14–24, Elba, Italy, 1993. COSIT '93, Springer-Verlag.
- [Ull95] Brygg Ullmer. Towards ‘Virtual Realty’: Explorations in spatializing web content. Submitted as a CHI '96 short paper, December 1995.
- [Yat66] Frances A. Yates. *The Art of Memory*. Pimlico, London, 1966.