Transitional Volatility in Web Navigation

Usability Metrics and User Behavior

David R. Danielson
Stanford University
Symbolic Systems Program
Master of Science Degree Thesis
June, 2002
Adviser: Terry Winograd

Chapter Overview

Table of Contents

1 Introduction

At its best, Web navigation design is supportive of and responsive to fundamental user behavior in hyperspace. At its worst, such design contributes to arguably the oldest and most devastating problem of Web navigation: the lost in hyperspace phenomenon. A Web user may commonly suffer from a basic problem: forgetfulness. She might forget where she is in relation to the remainder of the Web site, how she got there, where she was planning to go, and what information she has already seen, and moreover might be fully unaware of how to arrive at the information she has not seen, all commonly considered aspects of disorientation.

Navigation design does not just help her remember. Instead, it allows her to do what Vannevar Bush [Bush45] spoke of long before hypertext and the Web: to "reacquire the privilege of forgetting." Generally, Web navigators, like any other user of any other technology, are not primarily interested in remembering, or in learning the intricate details of a site; they are, rather, there to complete tasks [Carr+87]. Navigation design produces a level of confidence in a navigator, that she might set aside secondary concerns and less important information, but nonetheless that she "can find them again if they prove important" [Bush45].

The navigation process is one of forgetting and remembering, of taking some paths and ignoring others. At the heart of Web navigation behavior lies hyperlink following, an associative leap from one information chunk to another. In our discussion, we will be concerned with a crucial moment in the navigation process: the moment of transition. We will look at the sorts of ways we might characterize a movement from one Web page to another, and the transition's relationship to that complex interaction between user, task, and Web site.

1.1 Goals

At the highest level, the goal of this thesis is a better understanding of user behavior in within-site Web navigation. There are two primary sub-goals: (i) to empirically assess the concept of transitional volatility in Web navigation, which is generally concerned with (a) the navigational and content changes of the Web interface in page-to-page transitions, and (b) the user's ability to reorient herself to these changes; and (ii) to base the empirical investigation on a theoretical model of disorientation resulting from hyperlink transitions.

1.2 Motivations

The motivations for investigating intra-site Web navigation in general are quite simple:

- Designers can do something about the results. Knowledge regarding global navigation behavior benefits the design of Web browsers, but there is typically less a site designer can do about the inter-site behavior of a navigator than about her intra-site behavior.

- The area is ripe for theoretical inquiry and research. Much of our current understanding of Web behavior applies to global navigation [Catl+95, Taus+97a/b, Cock+01, McKe+01], and point the way to controlled studies of intra-site behavior.

- Navigation won't go away. Intra-site behavior in particular is much less likely to be predominantly made up of search than inter-site behavior, and will continue to be an essential skill on the Web.

- Many more transitions are within-site than between sites. Web navigation has a "patchy" character [Card+01].

- Studying navigation behavior on the Web is particularly challenging.

1.3 Transitional volatility

"Volatile" is a telling way to describe both the Web and a user's interaction with it. The Web itself and its individual sites rapidly change [Pitk98], and user interaction with it is rapidly interactive [Cock+01], causing navigators to encounter a wide range of design schemes, types of sites, and types of information. That is, we can speak of two sorts of ways the Web is volatile, one independent of user interaction, and one dependent on it. Another split, though, will be of more concern to our discussion: (i) an individual structure, such as a page or site, can change over time, and (ii) separate structures can be different in such a way that navigating amongst them produces a volatile experience.

Although not our focus in this discussion, the former is worth noting, since its volatility can interact with the extent to which navigators encounter widely disparate information and design. We can think of sites changing across time in terms of Newell's [Newe90] time scale of human action. Across the social band (days, weeks, and months), for example, news sites will alter their content, and any site is potentially subject to a major redesign. Across the rational band (minutes to hours), more specialized sites can be volatile, and news sites covering a big story might update content more frequently than is typical. And across the cognitive band (seconds), still more specialized sites might update content, such as with a stock site. Notice that these sorts of changes primarily deal with the site content. For an individual user, of course, the design scheme of a page might alter during a navigation session, for example based upon her navigation history within the site.

The sorts of changes above happen to an individual page or site over time. The concern here is with the differences between pages and sites a Web user navigates amongst. When she moves to a new site or sub-site, she is confronted with new navigation support and potentially different design schemes. These sorts of changes relate to the “flow” of a navigation session, and are indicative of an old concept in interface design and film production: the extent to which a scene is “visually turbulent” or has good “visual momentum” [Wood84]. There a number of reasons that the extent to which interface changes occur during a navigation session might be of special interest for many circumstances on the Web: (i) the page-to-page transitions are generally discrete and invoked by only one primary act on the part of the user, namely hyperlinking; (ii) the visual changes are often themselves discrete, but are affected by page loading time, which itself can produce variations on the rendering of a page; (iii) the Web is a wide open space in which users frequently encounter vastly different types of sites and vastly different design schemes; and (iv) the user is predisposed to look to the content of the site [Niel00], and thus the transitional changes on the Web will not necessarily be noticed in the same sorts of ways they are noticed with other interfaces; rather, the volatility of the transition (especially that of navigation support) may often go unnoticed until non-content interface objects are needed.

An empirical investigation of these sorts of changes, and the effects they have on user behavior and Web usability, might provide valuable considerations for Web information designers and information architects. This thesis aims to provide such considerations, with empirical support.

We cannot talk about page-to-page transitions in a vacuum, however. The user brings a background of previous engagement with the Web into each and every transition. In particular, we will say that she is more or less habituated within an information patch, and that attributes of the destination page are more or less predictable. Here is a simple way of looking at the interaction that will drive much of our discussion:

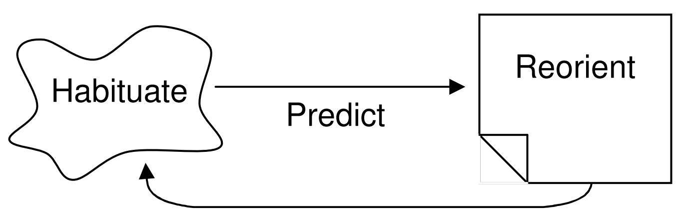

The user becomes habituated within the recent navigation patch. The user predicts content and navigation option changes in page-to-page transitions. The user reorients into the destination page of a transition. The destination page becomes part of the recent navigation patch, continuing the cycle.

One of the great problems in Web navigation is the “behind the door” problem: users are often unable to grasp what lies directly behind a hyperlink, let alone what lies further down the path of that hyperlink. Transitional volatility will be thought of largely in terms of its potential contribution to this navigation problem.

1.4 Roadmap

This thesis is composed of three main parts. Part I will provide a background discussion of intra-site Web navigation. In it, we will consider the major players in the Web navigation process: the hyperlink, navigational mechanism, navigational scheme, page, site, task, and user. We can think of the navigational constructs (all but task and user), as building up in aggregation:

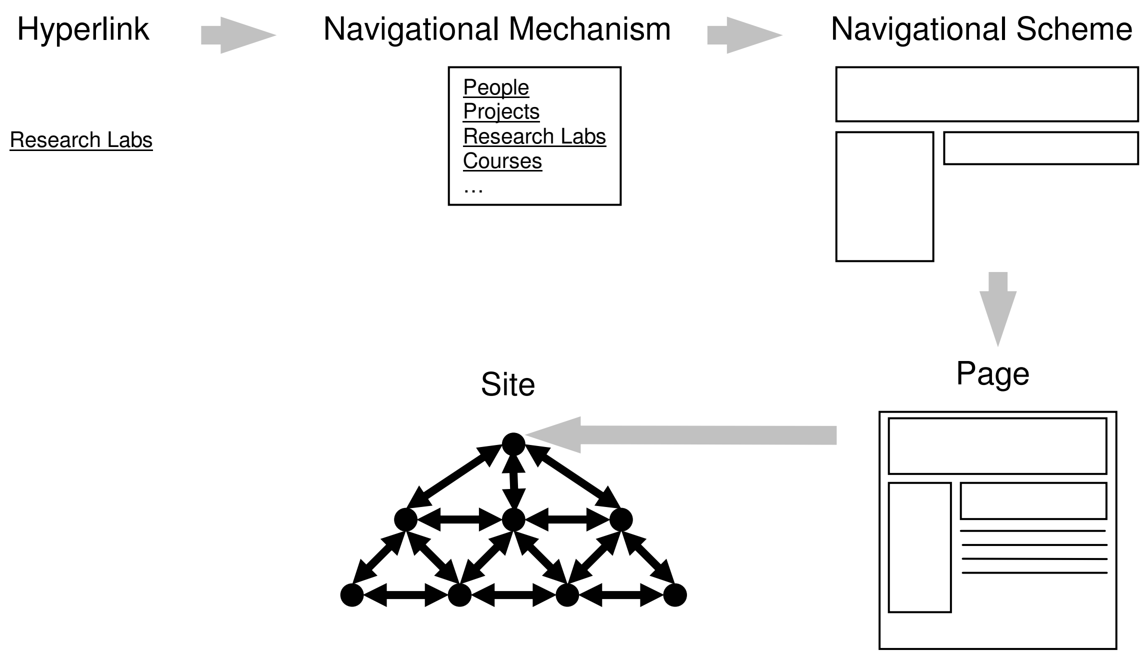

A navigational mechanism will be defined as a portion of screen real estate, with rules specifying which hyperlinks appear, and their display attributes. A navigational scheme is defined as a combination of one or more navigational mechanisms, with rules specifying which mechanisms appear, and their display attributes.

The goal of Part I is to map out the research space of the sorts of attributes we can consider, based upon empirical evidence and potential empirical investigation. Chapters 2-6 will consider these five constructs. In this thesis, we will be concerned with a particular kind of task, namely directed search, and will not investigate individual differences amongst users. For this reason, Chapters 7 and 8 will only briefly touch on some of the general attributes of the information need and the navigator herself.

Part II will discuss ways of measuring and assessing the complex interaction between user, task, and Web site, beginning with an overview of the research approach used in this thesis, namely click-stream analysis. We will pay special attention to a user's level of disorientation, her mental models of the site and of her navigation history, and to assessing a navigation session in terms of transitional volatility, habituation, and predictability.

Part III will present and discuss an empirical study of the effects of transitional volatility, based upon a proposed model of disorientation resulting from page-to-page transitions. The study will investigate influences on user behavior, disorientation, and perceptions of a Web site.

PART I: A FRAMEWORK FOR WITHIN-SITE WEB NAVIGATION

2 Hyperlink

The associative leap from one information chunk to another is at the heart of navigation in hyperspace. Bush called such leaps the "essential feature" of his conceptual Memex machine [Bush45], and Conklin [Conk87] later echoed the sentiment, calling them the "essential feature of hypertext." Web navigators make such associative movements predominantly via hyperlink following [Catl+95], and their behavior reinforces the hyperlink's central status in Web navigation interaction and design.

In the World Wide Web hypertext system [BernT+94], the linkable object in a document may be textual or graphical. When in text form, hyperlinks, as a standard, take on underlined and color-differentiated form, and therefore become visually distinct from the document's other text, not an atypical treatment for hypertext systems [Hala87]. The power of the hyperlink in Web navigation lies in its flexibility and simplicity. Web authors freely direct navigators to any other accessible Web document, of any document type, by authors who impose any design scheme and information structure of their choice, and all such movement is achieved through one method of action on the part of the user. In interesting ways, this simplicity gives rise to great complexity — in the way a navigator makes browsing and information-seeking decisions, in the way a navigator views the space and her available options, and, most notably with respect to our discussion in this thesis, in the way a navigator predicts characteristics of the hyperlink's resulting destination, and how she responds to that location given such limited predictive power at the source.

In the behind the door problem of hyperspace navigation, the hyperlink predominantly plays the role of the door. This chapter analyzes the deceptively simple hyperlink, paying special attention to those attributes which contribute to the navigator's ability to understand what's behind the door before she opens it. Given the hyperlink's central status, it will receive considerable attention here, to prepare us for the remainder of our discussion.

2.1 Typology

There is no generally agreed upon hyperlink taxonomy, but there are link types that typically surface in one form or another when one sets out to build such a taxonomy. A few dichotomies are common to hyperlink analyses — associative and structural, embedded and isolated, static and dynamic:

Associative/Structural

Associative: connect information chunks based solely upon content similarity and relevance.
Structural: connect information chunks based upon the site organization.

Embedded/Isolated

Embedded: visually and semantically included in the content of an information chunk in the Web interface. (Note an unfortunate overlap, common to classifications: "embedded" is often used, from the hypertext author's point of view, to refer to hyperlinks explicitly included in the HTML code [More+98]; this is equivalent to static hyperlinks defined below. The common usage provided here refers to the navigator's point of view, which we will prefer throughout our discussion.)
Isolated: visually and semantically separate from the content of any information chunk in the Web interface.

Static/Dynamic

Static: all attributes remain constant across navigation sessions and users; has no computed attributes.
Dynamic: has at least one computed attribute.

Notice the fundamental differences between these splits. Associative and structural hyperlinks define a relationship between the source and destination pages of the transition. Embedded and isolated hyperlinks define contextual, spatial, and design attributes. Static and dynamic hyperlinks define rules (or lack of rules) for the hyperlink's attributes, and indeed whether or not the navigational option appears to the user at all.

There are, of course, other hyperlink attributes we might tap into in building a useful typology. Park and Kim [Park+00a] make use of a site's connectivity to distinguish between basic links and add-on links. Basic links are a minimal set allowing a navigator to visit any page in a site, while add-on links are provided in addition to this basic set to improve navigation in the space.

DeRose [DeRo89] provides a taxonomy considering each of the three dichotomies above, giving special attention to the static/dynamic split at the top of the hyperlink hierarchy. Extensional links are hyperlinks which are stored individually, while intensional links define them functionally, for example by a search algorithm connecting destination content to a user query.

DeRose goes further in suggesting a tie between the static nature of a hyperlink and the predictability of a hypertext, and indeed provides glimpses of the importance of source-destination relationships. Extensional links are split according to this relationship; relational hyperlinks connect single conceptual units (equivalent to the term "information chunks" in our current discussion) together, while inclusive hyperlinks connect one conceptual unit to multiple conceptual units. This type of split will be discussed more in section 2.4, and under many structural conditions of a hypertext is closely tied to the organizational relationship of the source and destination pages.

2.1.1 Associative/Structural

The source-destination relationship of a hyperlink brings the issue of predictability more clearly into focus. Associative/structural classifications are in some sense primarily concerned with such relationships, and define hyperlinks in part by their relationship to the hypertext as a whole. Structural hyperlinks appear where they do in a hypertext as a result of the site organization; were the author to reconsider this organization, the placement of such links would alter, but the placement of associative links would not. Associative links rely exclusively on source-destination content similarity and an author-initiated notion of relevance, whether it be relevance of the hyperlink's destination to a likely user task, or relevance of the hyperlink's destination to its source. Note that it is necessary to say associative links rely *exclusively* upon these attributes; otherwise all hyperlinks would be associative. That is, structural hyperlinks do imply content similarity and relevance. Sibling nodes are not connected to one another simply as a matter of hierarchical convenience; that they are siblings in the first place ensures some level of relatedness and, therefore, association.

In real-world sites, associative hyperlinks are often given "see also" relevance to the current page [Niel00], typically because the source-destination relationship of the hyperlink is not already made explicit in the site hierarchy. Brown [Brow88] has made this distinction by referring to hierarchical and cross-referencing links. Hierarchical links are essentially structural, allowing the author to present information in sub-sites, sub-sites of sub-sites, and so on. Authors present cross-referencing links to allow for movement across the site's normal hierarchical boundaries, or to another site altogether.

For the purposes of this dichotomy, structural hyperlinks take on a slightly broader role than in some classifications. Otter and Johnson [Otte+00], for example, make a useful distinction between hyperlinks which are organizational and those which are indexical. The former maps more clearly to hierarchical structure, while the latter refers commonly to alphabetical listings of hyperlinks. In our current discussion, an alphabetical listing is simply one form of structure, although hierarchically organized sites are the primary focus of this thesis.

Perhaps the most notable difference between associative and structural hyperlinks lies in the predictive power they respectively bestow upon the navigator. The associative links of DeRose's taxonomy connect arbitrary pieces of information, and are strongly called "completely unpredictable" [DeRo89]. Associative, unstructured links have been compared to the *goto* statement in programming languages [Brow88] and its analogous predictability problems [Dijk68]. Some have called for an avoidance of associative links [Otte+00] primarily due to their contribution to the behind the door problem.

Predictability in hyperspace, as a concept, has been used somewhat loosely thus far, referring to a navigator's ability to anticipate attributes of a hyperlink, such as, most commonly, the information available at its destination. The concept will be discussed in more detail and context in Chapter 13.

2.1.2 Embedded/Isolated

The hyperlink's attributes in screen design on the Web have traditionally received less attention than organizational and functional attributes, but will receive due attention in this thesis. The embedded/isolated dichotomy refers specifically to the author's presentation of the link on the source page, and typically refers to the author's placement of the link in relation to surrounding body text [Vora+94]. Take, for example:

(i) Good Italian restaurants in Palo Alto
(ii) You might like to see the menus of good Italian restaurants in Palo Alto.
(iii) If you are interested in good Italian restaurants in Palo Alto, click here.

Hyperlinks embedded in body text make use of context to cue the navigator regarding the information available at the destination page, and can in theory promote the local coherence of a text [Vora+94]. Isolated hyperlinks have no such dependency. This "stand alone" property of isolated hyperlinks is crucial, as will be discussed in section 2.6. Recognize, however, that not every embedded hyperlink has a greater level of contextual dependency than every isolated hyperlink. Examples (i) and (ii) above demonstrate this; both links provide the same snippet of information, and (ii) could have easily been extended to include more. Example (iii), on the other hand, shows what is perhaps the most dependent textual link commonly used on the Web. Without context, it provides absolutely no information to the user other than its status as a clickable object, which is arguably immediately discernable by visual properties alone, at least for navigators with even small amounts of Web expertise.

2.1.3 Static/Dynamic

From a navigational perspective, perhaps the most notable difference between a search engine site and many of the information sites they lead users to is the reliance of the former on primarily dynamic hyperlinks, and the more common reliance of the latter on static links.

Dynamic hyperlinks are, in one way or another, generated by computation. The "results" page of a search engine may conceptually be the same page given any query, but its available hyperlinks are of course radically different for different information needs.

Dynamic hyperlinks span a greater range than the search engine example. Hyperlinks which change screen location, wording, or any other attribute on a page based on a user's navigation history, the time of day, or the color of the moon are all dynamic hyperlinks. Back button links are also dynamic. Regardless of the link attribute which changes, a hyperlink may only be referred to as static if none of its attributes are computed.

Our current discussion will be more concerned with static hyperlinks. Part of the motivation for this focus has been set forth in Chapter 1: within-site information-seeking is less dependent upon information retrieval. But note that the potential volatility associated with some types of dynamic hyperlinks (especially those based upon navigational history) may make for interesting future study.

2.2 Display attributes

The standard treatment of hyperlink text in the World Wide Web has been underlining and color differentiation, with blue and purple being used for previously uninvoked and invoked links. This specification benefits navigation by improving the scanning of navigation options on a page; visual search based upon color discrimination is known to be quite fast [Trie+80].

The hyperlink visual standard has had extraordinary implications for behavior and proposed authoring and design principles. Most notably, it has been widely argued and empirically validated that link titles are often scanned, read, and reacted to completely independent of the page content [Carm+92, Kapt93, Vora+94, Bach+97, vanW98]. This scan-browse behavior [Carm+92] has led to strong recommendations for focusing on hyperlink readability over global visual attributes [Kapt93], and for an emphasis on the stand alone quality of isolated hyperlinks.

The visual standard has become so powerful, in fact, that underlining non-hyperlink text on the Web can be a burden to navigators [Flem98, Niel00]. The standard is not only powerful, but noticeably simple. While hyperlink typing to more explicitly show link attributes and source-destination relationships has been proposed in many different forms, the basic tradeoff between distinction usefulness and type overload [Then+95] provides the basis for hyperlink visual simplicity.

The one color distinction hyperlinks do make use of has, as a matter of practice, shown to be a good choice, although there is little in the way of empirical evidence comparing other potential distinctions to the visited/unvisited split. The emphasis on navigation history and visited status predated, of course, empirical studies showing users revisiting pages at extraordinary rates [Taus+97a/b, Cock+01]. The effect of the initial hyperlink visual standard on user's revisitation rates is unknown. Had the initial visual standard focused on, say, intra- versus inter-site transitions, we might be faced today with drastically different user behavior on the Web.

Finally, consider that some hyperlinks are presented in such a way as to make visually explicit the structural relationship between the source and destination pages, most common to graphical maps on the Web. A lack of such graphical cues requires the navigator to discover the structure by making use of "larger patterns of proximity and typing" [Mars+93], and in fact such larger patterns are essential components of the contextual attributes of a hyperlink, and the predictability of the destination page.

2.2.1 Screen location

A push for isolated, stand alone hyperlinks has everything to do with a hyperlink's location on the page. One must consider at least two hyperlink attributes with respect to location: (i) whether the hyperlink is indeed isolated from the body text of the page, and (ii) whether the hyperlink is above or below the scroll line.

The first has implications on the time it may well take for a user to view the hyperlink after arrival to the destination page. Nielsen [Niel00] has argued for years, with empirical support, that navigators immediately look at the body text when arriving at a destination page, and consequently that a minimalist approach to navigation aids is wise. Of course, the matter may not be so simple. The user's information need has everything to do with this behavior. If she is looking for a navigational option she predicts will be available at the destination page, we may see different scanning behavior. The argument Nielsen makes speaks strongly to common user goals, namely, the predominance of looking for content at a destination as opposed to being able to predict a desirable navigation option at that destination. Were a designer to significantly improve the navigator's predictive power, navigation aids might become a more prominent information goal in and of themselves at the destination, a fundamental concept throughout our discussion.

Secondly, whether a hyperlink appears above or below the scroll line is not cut and dry, since navigators visit sites with vastly different screen resolutions. Nonetheless, should one conclude a hyperlink will likely appear below it, this characteristic could be crucial. The importance depends upon the likelihood of a navigator to scroll, which is not a completely harmonious subject in the literature [Chim+94, Bach+97, VanW98, Byrn+99b, Niel00].

Screen location can have more profound effects when a user has acquired some expertise with the particular site being navigated. Users can indeed in some cases invoke a link based upon its screen location without reading the hyperlink title [Kapt93]. The positional constancy and extent to which the user has relied on the link's screen location in past interaction [Teit+83] may well be a driving force behind such behavior.

2.2.2 Text and graphics

It is well known that pictures assist users in performing tasks involving problem solving and memory [Egid+88], and Web navigation often involves significant amounts of both. In many instances, graphics, rather than the standard underlined text, are used as hyperlink objects. Potential problems with such practice have not gone unnoticed [Flem98]; the power of the hyperlink visual standard and the widespread use of graphics on the Web can make it difficult to determine which graphics are hyperlinks and which are not. There is as of yet no empirical evidence pointing to semantic benefits of strictly graphical hyperlinks in Web navigation, whether they are worth a thousand words or not.

It is known, however, that graphics, used in conjunction with textual links, have positive effects. Egido and Patterson [Egid+88] found pictures to be useful navigation aids in menus during browsing tasks with a hierarchically structured pictorial database. They claim pictures play a disambiguating role for the verbal category labels, which may or may not describe the destination information well.

2.3 Contextual attributes

Hyperlinks do not live in vacuums, waiting to be clicked. We have already had a taste of the context of a hyperlink in section 2.1; the surrounding text can provide clues regarding both the destination page and the meaning of the underlined hyperlink text itself.

Particularly in hierarchical structures, even isolated hyperlinks can rely heavily upon contextual attributes in providing information about their destinations. For example, parent nodes commonly provide necessary context for understanding hyperlinks to their children pages. This thesis provides an example: a number of chapters contain a section called "Typology," but in order to fully understand the section's purpose, we must look to the parent chapter title.

There are a number of page features that feed into the context of a hyperlink, and most fall under two prominent areas. First, the widespread reliance upon lists as an organizational mechanism begs us to consider the list position of a hyperlink when included in such a structure. Secondly, a more general contextual feature to be considered is the extent to which a hyperlink must compete with other available navigational options.

2.3.1 List position

Lists have become the stock and trade organizational mechanism for Web navigation design. Designers rely on both vertical and horizontal hyperlink lists, as will be explored in Chapter 3. However, the common avoidance of requiring horizontal scrolling on the part of the navigator has led to heavy reliance on vertical lists.

Such design behavior has arguably made the wealth of knowledge attained prior to the Web on menu selection [NormK91] particularly applicable to analyses of navigation behavior. Indeed, this wealth of knowledge will be referred to throughout our discussion.

In some areas, however, menu scanning and selection behavior lack widespread agreement in the research literature, making it difficult to assume particular effects of position of a hyperlink on, for example, its likelihood of being invoked.

2.3.2 Navigational competition

The extent to which hyperlinks compete with one another has profound impact on navigation behavior in an information space [Piro+99, Park+00a/b, Dani01], which should come as no surprise. The specific way in which such competition affects behavior can vary according to the nature of the navigational mechanisms competing with one another, but that some attribute or another must stand out from the crowd for a hyperlink to be selected holds true in every circumstance. Quite commonly, the attention-grabbing attribute is the residue [Furn97], or information scent [Piro+99, Card+01]. The link's scent is in relation to a user's information need; when a link's attributes match the information need well, the link is highly competitive.

A few attributes contribute naturally to the competition. A link's status as embedded or isolated is one, since as we have mentioned the user is likely to immediately direct her attention to the page content. Additionally, wherever the link appears, we must consider some notion of screen density [Dark+93]. On the web, the concept can be difficult to consider and employ in a practical way. We know that users make heavy use of scan-browse information-seeking strategies, often skipping non-hyperlink text and directing search focus to the visual hyperlink standard. The potential implication is clear: no matter the amount of surrounding text, embedded hyperlinks should be thought of as typically existing in a much less dense screen environment. Isolated hyperlinks, on the other hand, are typically part of a navigational mechanism, and therefore more rigorously compete with other navigational options. Regardless, screen density becomes a more important issue when one considers the rapidly interactive nature of Web navigation; the typically small amount of time a navigator spends on any given Web page pushes the importance of clearly visible hyperlinks [Cock+01]. Density may also play a role in the essential danger associated with including too many of Park and Kim's [Park+00a] add on links — the navigational overload of a page [Edwa+89].

A fundamental problem in analyzing navigational competition is in attempting to predict the navigator's mindset based upon the screen location of her eye gaze. We might hypothesize that focus on a navigational mechanism likely indicates a navigational mindset, while focus on page content more likely indicates an information extraction mindset, in which case embedded hyperlinks are at a competitive disadvantage, since the navigator would be less ready to move on when passing them over. Research in this area is necessary to better understand the meaning of screen location focus, and its effects on navigational competition.

2.4 Relationship

Hyperlinks, by nature, define and indicate a relationship between source and destination information chunks in a hyperspace. Moreover, they represent a hypothesis by the hypertext author that the navigator will be interested in the destination page [Pitk+97], given her current location at the source. The relationship can be general, such as "the destination provides additional information about the source," as is the case with many of DeRose's annotational links [DeRo89], or specific, such as "the destination page shows the menu of the selected restaurant on the source page." Specific relationships typically must be inferred by the navigator in hypertext systems. Typed links can make these semantic connections explicit [Thür+91], but as has been mentioned, the simplicity of Web hyperlinks avoids such specificity; they indicate a relationship, without the semantics.

Thüring, Haake, and Hannemann [Thür+91] distinguish between three levels of specificity in hyperlink source-destination relationships. At Level 1, the link carries no label, representing the most unspecific connection between the source and destination pages. At Level 2, a global source-destination relationship is defined, such as "the source is discussed by the destination." At Level 3, a more refined semantic relationship is indicated, such as our example of the restaurant menu at the destination page.

As Web information spaces are often hierarchically structured, hyperlinks commonly define a general, structural source-destination relationship, giving rise to three hyperlink attributes, familial relationship, direction, and distance. The hierarchical relationships themselves can be associated with the specificity of Level 2 or 3 in the Thüring, Haake, and Hannemann system. Often, the site hierarchy itself may be an IS-A or PART-OF hierarchy, such as with a site presenting the animal kingdom in the traditional IS-A organization, with a page devoted to each animal type.

2.4.1 Target

Despite the strength and pervasiveness of the page metaphor, the source-destination relationship of a hyperlink is not necessarily as simple as the relationship between its source and destination pages. The specific information target of the hyperlink gives rise to a complicated information-seeking process in Web navigation. The hyperlink may be intended to point to an entire site or sub-site, a page in the space, or a small portion of a page [Hala87, Furu+90], as in Figure 2.1.

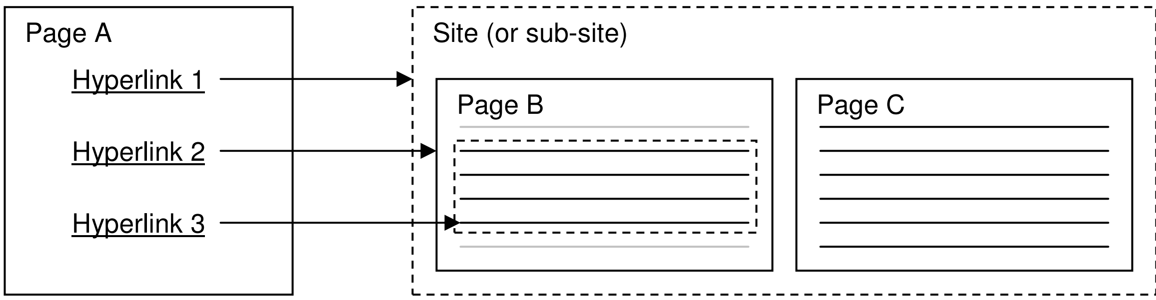

Figure 2.1: Hyperlink target specificity. Entire site or sub-site (Hyperlink 1), entire node (Hyperlink 2), or specific information chunk within node (Hyperlink 3).

Some obvious questions come up. To what extent does the information chunk Hyperlink 3 points to "pop out" at the destination page? That is, how quickly and easily can the navigator locate it given her information need? Is the information target above or below the scroll line? Hyperlinks pointing to specific information chunks at a destination page pose special problems for information seekers, who now must deal with other competing and peripheral information upon arrival. The navigator's ability to match her information need at the time of the hyperlink click to the specific information chunk at the destination has much to do with the site designer's use of meaningful text aggregates [Mars+93], for example in heading the target information chunk with the same text snippet as the proximal hyperlink. A common tactic is to direct the navigator to the beginning of the target information chunk using an anchor, although the effect of such practice on disorientation in hyperspace is relatively unexplored.

Notice that hyperlinks with specific information targets are likely relational, connecting to one information chunk. When hyperlinks point to full-page targets, we must investigate the attributes of the destination page in order to fully understand the hyperlink at the source. Hyperlinks may point to, for example, an index or an overview page, meant primarily to aggregate hyperlinks for site organization, in which case the link is inclusive, or may point to a content-centric page, typically low in the site hierarchy, in which case the link is relational. The singularity or plurality of the information target, of course, is not necessarily easy to determine.

2.5 Information scent

Hyperlinks may be viewed as serving as proximal cues to information chunks and pages in the space. From the point of view of the hyperlink, there is some target (or set of targets) for which it provides scent. From the point of view of the information target, there is some set of information chunks and hyperlinks spread out in the site for which it has residue. Information scent is a primary concept in the information foraging theory proposed by Pirolli and Card [Piro+95, Piro+99]. Residue, a similar concept from a slightly different perspective, was introduced by Furnas [Furn97]. Both speak to the proximal cues Web navigators are flooded with during information-seeking. At an intuitive level, given some information need, a Web navigator makes navigational decisions based upon proximal cues she believes will likely lead her to the desired information target, called scent-following [Card+01].

2.5.1 Relevance

Following scent in Web navigation carries with it a notion of relevance between the hyperlinks and information chunks that provide the scent and particular information need or interests of the navigator. In some sense, the relevance of a hyperlink to the navigator's needs and interests matters the most. She arrives at the site and moves around within it in the context of some sort of explicit or implicit goal or set of goals, whether it is to find a specific piece of information, to read something interesting, or to simply waste time. Hyperlinks are relevant to some degree in relation to those goals [Piro+95]. They indicate to the navigator whether the destination is likely (or unlikely) to lead to desired information, something interesting, or a big waste of time, and under considerations of relevance the hyperlink titles are typically referred to as the surrogates [Mizz97]. The quality of that surrogate in closely matching its described document relies in part upon the expressiveness of the hyperlink, described in the next section. Mizzaro [Mizz97] explains that, with respect to information retrieval, the quality of a surrogate title may or may not increase with length.

2.6 Expressiveness

From one perspective, a hyperlink may be viewed as a symbolic representation of its target information chunk, whether that be a Web site, a Web page, or a paragraph of content. A hyperlink's expressiveness answers a simple question: How much does the symbolic representation at the source tell the navigator about the destination target it represents? Given such a view, we may first talk about those features of the target which the hyperlink makes *explicit*. Johnson-Laird [John88] illustrates symbolic explicitness with the "prisoner's tally." For example:

(i) II
(ii) IIII

Which is bigger? You can quickly and easily determine that the number represented by (ii) is larger than that represented by (i), simply by the length of the tally *and nothing more*. In one way, we have seen that a hyperlink behaves for a navigator in a similar (but not equivalent) way, making clear the target's status as previously visited or unvisited by quickly recognizable color coding. Some hyperlink attributes, on the other hand, may not necessarily tell us much about the target, such as the length of the text snippet; some hyperlinks are just wordier than others. (Although, in context, the hypertext author could conceivably use length as an explicit cue.)

Recognize, however, that the explicitness of the length of a prisoner's tally arises naturally out of the system of which it is a part; one could not cause the tally length to stop providing numerical size information without altering the tally system altogether. Hyperlink color coding, on the other hand, is merely a convention, but has become so powerful a convention that one could argue it *acts* like an explicit cue. The majority of hyperlink attributes providing information regarding the target do not make target attributes explicit. Rather, the navigator must use the context of the site (and the Web) as a whole to recognize and benefit from design conventions. As a simple example, the screen location of a hyperlink could be used to indicate the specificity of the target — links towards the top of the page pointing to sites and sub-sites, links midway down pointing to single pages, and links near the bottom pointing to specific information chunks on a page. Any such author convention does not make target attributes explicit, but rather adds more generally to the expressiveness of the hyperlink. An obvious issue now arises, regarding the extent to which conventions such as this screen location example provide natural mappings (in this case, higher up on the screen mapping to greater generality of the target) between the hyperlink and target attributes. Moreover, we must consider the navigator's ability to recognize such mappings given, potentially, only a short time at the site using the mappings.

2.6.1 Semantics

Hyperlinks are generally distinguished from one another by their content, whether it be a snippet of text or graphical content. These snippets are the primary sources of immediate empirical evidence upon which users make navigational and information-seeking decisions, and thus Web navigation is, by nature, largely semantically driven. Hyperlink titles can be as expressive (or unexpressive) as the sentences we utter, and describe the destinations they lead to, helping to reduce disorientation [BernM00]. That is, they are the proximal cues to the distal content [Chi+01], as we discussed in section 2.5.

There are three major problems a Web author faces in presenting hyperlink titles. First, she could fail to describe the destination well, or label a link attribute inappropriately [Thür+91]. Or, she may succeed in describing the destination but fail in describing a larger, sometimes more important picture: what Furnas [Furn97] refers to as the whole-to-set of the hyperlink, which we should associate with links targeting entire sites or sub-sites. For our present discussion Furnas' analysis is crucial: "This is a problem in many hypertext systems, including the WWW: their link-labels indicate adjacent nodes and do not tell what else lies beyond" [Furn97]. Furnas acknowledges that providing such cues can be difficult in practice, but as we will see throughout this thesis, there are a number of contextual tactics designers can employ to help navigators understand the entire path or subspace off in the direction of a particular hyperlink, and a number of behavioral conditions under which a user can better make such predictions. The pervasiveness of insufficient link semantics have several negative information-seeking consequences. Otter and Johnson [Otte+00] point out that navigators are likely to rely upon more general (and often less effective) search strategies when hyperlinks provide little semantic help, and that hyperlinks must therefore prime the user for links further along a navigational route, something we will discuss in Chapter 13. An obvious problem arises when the author writes a poor hyperlink title: the navigator simply misinterprets the author's intention [Gray90].

Secondly, and more worrisome, misunderstandings about a destination based upon hyperlink titles are quite likely impossible to avoid altogether. Hyperlink descriptions are susceptible to the vocabulary problem in human-system communication [Furn+87]: people tend to show great variability in their word and term choices for describing a concept.

Thirdly, hyperlink title writing falls under one of the most basic tradeoffs in all of human-computer interaction design, pointed out by Norman [NormD83]: "Factors that increase informativeness tend to decrease the amount of available workspace and system responsiveness." Very expressive hyperlinks are likely to take up too much space, and potentially require not only system resources, as Norman points out, but great cognitive resources on the part of the navigator as well. The title writer faces the basic problem of attempting to make hyperlinks short and sweet, and the resulting lack of expressiveness tends to lead to occasional hyperlinking that is followed by an immediate Back button click because the user is surprised (and dismayed) by unexpected information at the destination.

Systematic attempts to avoid unnecessary link firings and misunderstandings of the sort described above have been called link destination announcements by Germán and Cowan [Germ+00], and these can be achieved by semantic consistencies in a Web space. But these announcements need not be explicit. Consider the hyperlink Terry Winograd. You might expect such a link to direct you to a home page, or perhaps to open your email editor, given past experience with linkable proper names on the Web in general. If you are visiting a bookstore site, you might reasonably expect it to lead to a list of books on sale. Moreover, you may allow precedence within the current Web space to guide your predictions, for example if other proper name links within the site have all led to home pages. Notice that the hyperlink title tells you little about the destination: only that it has something to do with a particular person. Because of your previous interaction with similar linguistic structures, the hyperlink becomes more expressive, but only because of a certain level of consistency throughout the site. The extent to which transitional volatilities are homogenous, in one way or another, will be critical throughout our discussion:

> *Linguistic Homogeneity of Volatility:* the extent to which hyperlinks (across a site) with similar linguistic structures result in similarly volatile transitions.

The notion of consistency turns out to be rather tricky, and we will return to it a number of times. The rules a Web author employs can be consistent in all sorts of ways in a technical sense, but utterly inconsistent (and unpredictable) from the point of view of the navigator. A proper name might lead to home pages under particular circumstances and open an email editor under others. A largely unexplored empirical question regards the extent to which navigators are able to learn and benefit from such author-initiated hyperlink rules, and generally how complex the hyperlink rules can afford to be before the Web interaction tends to break down.

Finally, the semantic expressiveness of a hyperlink title can bestow important predictive powers upon a navigator regarding the design of the destination page. Consider:

(i) See my page for good Italian restaurants in Palo Alto.
(ii) See my list of good Italian restaurants in Palo Alto.

Presenting (i) provides no real indication of the design structure. Sometimes this can be for good reason. Such structures can change, so if the Web designer keeps the residue of a page (such as hyperlink titles leading to it) ambiguous with respect to destination design, she is free to change the design structures without producing a wave of misleading scent throughout the site, which she would now need to fix. Authoring concerns like this are largely neglected, and are indicative of one of many chicken and egg problems on the Web; we simply do not know if the lack of authoring support has caused ambiguous link titles with respect to design, or vice versa (or, alternatively, if the problem is perceived to not be worth the trouble). In any case, there are clear advantages to presenting (ii). Suddenly the navigator knows what to look for as the page is loading — a particular type of design structure she has undoubtedly encountered many times before. How will this affect her behavior at the destination? Will she be more likely to neglect any and all content isolated from the list structure she is primed to see? The complex ways in which Web pages are rendered make these difficult questions to answer. We will see in Chapter 10 that different empirical approaches can yield different types of insights on this matter, but the effects of subtle semantic decisions such as this on the Web remain an open area for interesting research.

2.7 Identity

There is a final question left unasked thus far, but important throughout our discussion: What is the scope of a single hyperlink's identity? From a technical perspective, the question is admittedly a silly one, but from behavioral and design perspectives, it is not. When do navigators perceive of and treat hyperlinks across Web "locations" and time as a single entity, and under what design conditions may we assume and benefit from hyperlink identity?

The question naturally extends to navigational mechanisms and schemes, pages, sub-sites, and sites, and becomes more complicated with these conceptually larger and more volatile structures. If we agree that behavioral measures are to be the benchmark by which we determine identity and coherence of navigational structures on the Web, identity is not easy to establish; we have at current no reason to believe identification of such structures follows the identification of objects in the physical world in any way. In fact, the prospects may be even more grim. It should not be taken for granted that a navigator *ever* perceives or treats two technically distinct navigational structures as identical.

3 Navigational Mechanism

The power of the hyperlink may partly rely upon its flexibility, but the power of an information space often relies upon its organizational structure. Free reign in associative linking has arguably played a major role in the lost in hyperspace phenomenon. Without organizational structure, a navigator quickly and easily forgets the context of her movement through a space.

Navigational mechanisms on the Web provide the first level of aggregation in site design. Recall our definition from Chapter 1: a navigational mechanism is a portion of screen real estate, with rules specifying which hyperlinks appear, and their display attributes. Mechanisms range from ambiguous “see also” associative hyperlink lists to full-fledged site maps. One of the great challenges in navigation design is to cause the navigator to view such mechanisms as single conceptual entities, with their own coherent navigational purposes.

In this chapter we will discuss the issues surrounding navigational mechanisms, the most basic organizational structures in navigation design, looking at some of the common mechanisms used on the Web and their implications for page-to-page transitions.

3.1 Typology

A navigational mechanism taxonomy parallels much of the hyperlink taxonomy we saw in Chapter 2. Indeed, one might like to classify mechanisms by the hyperlinks they aggregate. Common mechanisms present purely associative hyperlinks, often grouped with “see also” or “related pages” labels, to name a couple. The embedded/isolated split applies less commonly to navigational mechanisms than to singleton hyperlinks, primarily as a matter of design practice [Haas+98], but short hyperlink lists do appear embedded in page content.

Navigational mechanisms, by definition, typically have a dynamic nature, applying rules to determine hyperlink attributes, although, as we will see in section 3.2, we could take a slightly different approach to the static/dynamic split in classifying mechanisms in a useful way. Dynamic navigational mechanisms that present temporal information, such as history lists, have received considerable attention [Taus+97a/b]. Park and Kim [Park+00a] distinguished between structural and temporal navigation aids, essentially comparing examples of static and dynamic mechanisms, if we treat all structural groupings as static (given a structure that is itself static).

3.1.1 Local/global

Another dichotomy tends to surface in analyzing groups of navigation aids, namely a local/global split [Park+00a/b]. Such a split often refers to the hierarchical levels of the hyperlinks' destination pages. A set of hyperlinks pointing to the top-level nodes in a site make up a common global navigational mechanism, for example. Mechanisms with global context are more likely to remain static than those with local context, which typically use rules with greater specificity.

Bernstein [BernM98] refers to the local/global split in terms of a given page's neighborhood, suggesting that, "A Neighborhood establishes an association among nodes through proximity, shared ornament, or common navigational landmarks." Particular attributes of a navigational mechanism may identify the mechanism as providing local support.

3.1.2 Site/Browser

User behavior provides us with another important dichotomy. Catledge and Pitkow [Catl+95] reported that other than the browser's Back button, navigators rarely interacted with browser-supported mechanisms. Navigational mechanisms appearing in the user's Web browser are fundamentally different from those appearing within a site. First, browser-supported mechanisms, such as the Uniform Resource Locator (URL) text field and the navigator's history list, must be generalized across any encountered Web site, and so, with the exception and the Back and Forward buttons which uniquely support recency, browsers are typically used for inter-site navigation [Bach+97].

We will return to the Back button in section 3.7, and the lack of use of other browser-supported mechanisms has been an area of interest for some time. Byrne, John, Wehrle, and Crow [Byrn+99b] suggest that URL typing requires too much time and cognitive effort: "URLs appear to be particularly difficult for users to type because of the unusual punctuation and preponderance of nonwords." Thus, navigators are prone to error in attempting to type URLs, particularly for pages deep in a site hierarchy. Additionally, in the Catledge and Pitkow [Catl+95] study, "users typically did not know the location of documents a priori," which may contribute to infrequent URL typing. Perhaps most importantly, URL entry typically requires recall rather than recognition memory (although newer browsers attempt to address this problem with a URL history list). Lastly, URLs are possibly closer to the underlying structure of a hyperspace than many users care to think about.

Lack of history list inspection presents another interesting mystery in navigation behavior. Tauscher and Greenberg [Taus+97a/b] suggest that poor design of browser history mechanisms explains the lack of use. However, the Back button may provide all the history support navigators need if they "rarely back up more than a page or two at a time" [Byrn+99b].

Notice that because browser-supported mechanisms typically must generalize across sites, they are predominantly temporal in nature, while within-site mechanisms may be either structural or temporal. We will discuss temporal navigation behavior in Chapter 10, and see that a big question for browser and site design remains: Just how much temporal support would or even could theoretically be useful in Web navigation, and more specifically in within-site Web navigation?

3.1.3 Size

Now that aggregation of navigational support has entered the picture, we must consider the size of the aggregated navigational mechanism. Size is often a matter of the number of hyperlinks presented in the mechanism, although this assumes some sort of uniform hyperlink display within that mechanism, which is common. Number of hyperlinks is particularly important with respect to site overviews. Chimera and Shneiderman [Chim+94] distinguish between large, medium, and small (vertical) overviews in terms of the required screens to fit all overview items:

Small: all items fit on one display screen

Medium: between 2 and 15 screens are needed to display all items

Large: more than 15 screens are needed to display all items

In Web interface terms, the size of an overview depends in large part on the necessity to scroll in order to view all items. Darken and Sibert [Dark+93] provide a similar perspective on virtual worlds more generally with definitions for “small” and “large” worlds:

Small World: a world in which all or most of the world can be seen from a single viewpoint such that important differences can be discerned

Large World: a world in which there is no vantage point from which the entire world can be seen in detail

3.2 Display attributes

The presentation of a navigational mechanism, even deceptively simple hyperlink lists, involves a number of design decisions, many of which carry over from menu design [NormK91, Shne98, Zaph+00]. These decisions can be important in shaping users’ navigational behavior and perceptions of a Web site.

3.2.1 Screen location

As with individual hyperlinks, the screen location of a navigational mechanism is a crucial display attribute. As usual, the scroll line is a potentially crucial factor; reducing or eliminating the necessity to scroll to hyperlink locations is arguably critical for usable site design [Bach+97].

An investigation of frame layout effects by van Schaik and Ling [vanS+01] revealed that navigation frames presented at the top or left of a page improve visual search performance. The advantage of left-column hyperlink frames coincides with the practice of "yellow fever" [Niel00], in which navigation options are commonly presented in a column along the left side of a Web page, usually with a background color distinct from that of the page content. Cognizant of this common practice, van Schaik and Ling tested but found no differences between new Web users and those more familiar with the Web, and thus claim that exposure to the practice did not influence performance. They further acknowledge that the results appear to coincide with Western reading patterns (top-left to bottom-right), but that cross-cultural research is required to fully understand the effects of frame positioning.

3.2.2 Hyperlink presentation

Navigational mechanisms predominantly appear in the form of textual hyperlink lists. Moran [Mora98] refers to "visual impact" as a fundamental presentation concern: a navigator needs to immediately understand that she is looking at a list. Moran points out that one of the most prominent tactics used in achieving this visual impact is presenting lists isolated from other content, perhaps a cause of Web designers' typically strong reliance on isolated hyperlinks reported by Haas and Grams [Haas+98].

For hyperlink lists, Web designers have three other primary presentation issues: (i) promoting easy scanning, (ii) making particular relationships between items within the list explicit when appropriate, and (iii) determining the order of the hyperlinks in the list.

Many of the particulars of hyperlink list presentation can be and often are borrowed from menu design, for example in flushing the most important words in a list item to the far left to be supportive of scanning behavior [Shne98].

Traditional overview presentations demonstrate item relationships primarily via indenting, or, with more graphical presentations, simple lines connecting related items; first and foremost, however, Moran [Mora98] points out that list relationships are denoted by parallelism, or the connection of all items in the list by way of some title. Moran further suggests the importance of consistency in the grammatical structures used for list items. The point is intimately related to the extent to which list items result in homogenous transitions, which we will discuss more in section 3.6.

Finally, hyperlink order is a potentially important presentation attribute. As Turo and Johnson [Turo+92] point out, hyperlink order may provide useful information about a destination page's importance, amongst other attributes. Two concerns regarding the navigator herself will better inform us about the potential importance of hyperlink order: (i) her propensity to look for and understand the sorts of implicit messages about the destination page by the order, and (ii) her list scanning behavior, which we will come back to in Chapter 5.

3.2.3 Mechanism rules

The crux of the navigational mechanism is the set of rules it lives by. For example, the Back button, which we will discuss in section 3.7, is typically stack-based. That is, the rules for which page the Back button itself links to, and, in the case of many browsers, the recent history list in the button's pull-down menu, are stack-based. Each visit adds the new page to the stack, until the Back button is invoked and a page is popped off the stack (and pushed on to the Forward button's stack, which will not be a large concern due to its infrequent usage).

Rules for navigational mechanisms, as one might guess by our typology discussion in section 3.1, can be more or less complex and dynamic, and more or less based on the user's navigation history or the site structure. Structural mechanisms use the site hierarchy as the basis for their presentation rules, for example in following a rule to always display links to the top-level category pages in the site horizontally across the top of the page. The extent to which these mechanism rules present different hyperlinks from page to page can vary considerably. All pages, for example, may have a different set of children (and do if we assume no repeating nodes), and so a structural child list mechanism will present a different set of hyperlinks at each page in the site. Moreover, a history-based mechanism may present a different set of hyperlinks on the same page at different points in time, and so the level of dynamism of a mechanism cuts across two spectrums: (i) the extent to which the mechanism will change due to page-to-page transitions (between pages), and (ii) the extent to which the mechanism will change within a single site location (within pages). An important and open area in Web navigation continues to be the extent to which navigators understand mechanism dynamism, and how such dynamism affects their behavior and attitudes.

3.3 Contextual attributes

Just as the level of dynamism of a mechanism cuts across the inter- and intra-page spectrums, the context of that mechanism, too, must be viewed from these two perspectives. Mechanisms, by our definition, exist across a hyperspace, relying on screen location to achieve identity. The intra-page context of a mechanism, however, poses more immediate usability concerns, competing with other interface objects.

3.3.1 Competition/coordination

There is a cost associated with the presentation of each additional hyperlink; like hyperlinks, mechanisms must compete and coordinate with other mechanisms, and the dangers of navigational overload apply [Edwa+89, LyncP+97]. Little is known about the specific ways mechanisms compete. Kim and Yoo [KimJ+98], for example, looked specifically at cyber mall tasks and reported that neighborhood and top-level links, when presented together, appeared to increase the convenience of searching, browsing, and satisfaction.

3.4 Relationship

Mechanisms define an aggregated, and sometimes complex, relationship between a set of destination pages to a source page. Moreover, they typically can, either implicitly or explicitly, indicate relationships amongst the destination pages, starting with Moran's [Mora98] parallelism. Note the possible beginnings for indicating navigation beyond adjacent nodes: by demonstrating a relationship between two hyperlinks at the source page, there is potential for divulging navigational scheme attributes at the destination of a hyperlink. It is not clear that in practice such demonstrations of hyperlink relationships allow for more sophisticated predictions about the navigational schemes at destination pages, but they are a theoretical source of such information and would make for interesting investigation.

3.4.1 Guided tours

The extent to which a Web space is essentially linear will be an important issue throughout our discussion, and one of the most common approaches to presenting linear hypertext is that of the guided tour, in which the Web information designer usually forces (but at least implies) a user to take a particular navigation path through the space [Mars+89]. Guided tours are traditionally accompanied by "previous" and "next" progression hyperlinks, which when presented in their typically ambiguous fashion can be mistaken for Back and Forward navigation [Bach+97].

Note here that in some sense every navigational mechanism may include some implicit element of progression or guidance, since the ordering of the hyperlinks may imply a progression (i.e. from left to right for horizontal lists, or from top to bottom for vertical lists).

3.5 Hierarchical

The study described in Chapter 15 of this thesis deals specifically with navigational schemes for hierarchically organized Web spaces, and so it will be useful to discuss in this section some of the primary hierarchical mechanisms. Hierarchical mechanisms are often graphically distinguished from associative and cross-referencing hyperlinks through what are often called navigation side bars [Øest99], and these side bars tend to present groupings based on particular kinds of hierarchical relationships, such as a listing of the siblings of the current page.

3.5.1 Top-level

Top-level is a loaded term when dealing with hyperspace. As in real space, there are no inherent directions on the Web; they are instead imposed by the information designers. Commonly, there are a small number of top-level pages (recommended in some publications to be between five and nine, often based upon the magic number proposed by Miller [Mill56]), from which a path to all other nodes exists. There are special dangers associated with presenting too many top-level links, such as the potential for increased feelings of lostness [Lars+98], especially given that for primarily hierarchical sites the initial path decision can be critical to information-seeking success [NormK91]. Top-level links are generally to the pages without ancestors in the hierarchy (although they are sometimes children of a single home page), and are an example set of basics links, defined as “a set of minimum connections that enables users to visit any node in a Web system” [Park+00b]. It is a common practice to include top-level links consistently throughout a site, and the links may often be found in the left column of the page or horizontally presented across the top of the page, or across the bottom of the page [Haas+98].

In addition to being loaded, top-level is often a relative term. It can refer to those pages mentioned above who generally have no ancestors, or, given some present location of the user, it may refer to the top-most hierarchical pages in a sub-site, or to the siblings of the parent of the current node (its aunts and uncles), or, indeed, anything in between. Commonly one does not find more global context than one or two relative top-level lists presented on a page.

3.5.2 Local context

The local context of a node within a hierarchy generally includes the node’s siblings, the node’s children, or both, and along with top-level and cross-referencing links often provides a simple fisheye view of a Web space, which we will briefly discuss in the next chapter. Local context is often particularly important for breadth-first navigation, although it supports such a strategy incompletely. Specifically, a level of a hierarchy includes several groups of siblings (for example, if there are five top-level nodes, the second level of the site consists of five sets of siblings). Local views do not commonly provide hyperlinks to cousin nodes, and thus moving through a level of the hierarchy would require visiting a top-level page as an intermediate step in arriving at a desired set of cousin nodes. This has likely contributed to some researchers suggesting that current design practices promote depth-first navigation, while ignoring the potential of breadth-first search. Newfield, Sethi, and Ryall [Newf+98], for example, claim, “breadth-first navigation permits a user to maintain more context as they search the Web, especially that relating to prior intention.” Additionally, breadth-emphasizing design results in the user being “constantly reminded of the full scope of services available on the site,” which is “particularly useful for users who do not enter at the home page but go directly to a page deep within the site” [Niel00]. The tradeoffs make local context presentation tricky, however. As Modjeska and Marsh [Modj+97] suggest, horizontal cross-referencing hyperlinks can often increase efficiency at the cost of orientation.

3.5.3 Breadcrumb lists

Both top-level and local context mechanisms generally provide explicitly structural hyperlinks, revealing hierarchical organization. Breadcrumb lists serve a similar purpose, but almost always do so in a strictly depth-emphasizing fashion.

The name of course is borrowed from navigation in physical spaces. Much like the physical breadcrumbs dropped in the woods to find one's way back home, hyperspace breadcrumbs provide context information commonly regarding where a navigator has already been, often presented in the form "Page A > Page B > Page C" where Page A is the parent of Page B, and Page B is the parent of Page C, which is the current node and therefore often not presented as linkable text.

Breadcrumb lists in hyperspace are a way of showing a navigator how she either did or might have arrived at the current location, and so typically represent a beaten path navigation aid [Paru89] in that the available options are likely to lead to pages already visited. The likely path to the location, most often the depth-emphasizing path from top-level to the current node via a series of child nodes, is presented. However, in some sense the physical metaphor breaks down, since if the user arrives at the node through the "side door" (does not take the standard depth-emphasizing route), then there will be no one there to drop the breadcrumbs in the intermediate steps, save the information designer herself. Breadcrumb lists are more appropriately thought of as structural contextual aids by which the information designer deduces the likely or initially envisioned path to the current location.

Nielsen [Niel00] has contended that "breadcrumbs are useful only for hierarchical information architectures because they require nested levels of progressively smaller subsites." One can imagine breadcrumbs of all sorts, though. Unlike physical space, discrete movements in hyperspace need not be to adjacent locations within the navigation structure, and so a breadcrumb could conceivably be "dropped" at any possible sequence of nodes the information designer deems a probable or logical navigation path. Guided tours, as we discussed in section 3.4, are a nice example of where breadcrumbs might otherwise be used.

3.6 Homogeneity

For the purposes of our present discussion, the extent to which a navigational mechanism's hyperlinks are homogenous is an important attribute. Our discussion of Moran's [Mor98] analyses in section 3.2 touched on the issue. Sets of hyperlinks can share more or less homogenous linguistic structures in their text snippets, and can be more or less homogenous in presentation. The homogenous attribute that we will primarily be concerned with is that of the sorts of interface changes occurring from source to destination as a result of following the hyperlinks in a particular mechanism. Take for example the following mechanism, likely to be found on a University department Web site, with a description of the hyperlink destinations on the right:

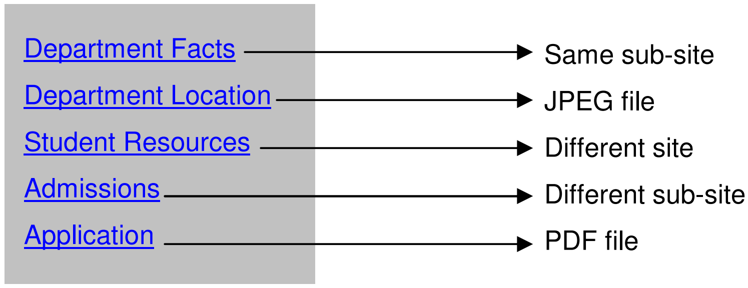

The hyperlinks in this single mechanism point to a variety of file types; the department location hyperlink might point to a JPEG map, and the application link might point to a PDF file to be printed and completed. Furthermore, those pointing to HTML documents differ in the relationships they present between the source and destination pages, with some leading to a destination within the same sub-site and others to different sub-sites or different sites altogether. The differences between the sorts of transitions hyperlinks within a mechanism result in, then, can be subtle, especially when each of the hyperlinks point to pages within a Web space using unifying design attributes and similar schemes. The general attribute we will concern ourselves with is the homogeneity of a mechanism's volatility transitions:

> *Mechanism Homogeneity of Volatility:* the extent to which hyperlinks within a navigational mechanism result in similarly volatile transitions.

Movements in a Web space between pages with similar structural relationships often have similarly volatile transitions, since pages within a local neighborhood typically share navigation support and other design attributes, and so whether a mechanism is explicitly one of the structural lists mentioned in section 3.5 is intimately related to its homogeneity of volatility.

3.7 Back button

Studies by Catledge and Pitkow [Catl+95] and Tauscher and Greenberg [Taus+97a/b] have made it clear that at current the Back button is the most critical browser-supported mechanism. Greenberg and Cockburn [Gree+99] point out that the Back button allows for rapid return to recently visited pages, and can be used with a simple "click until the desired page is recognized strategy." That is, Back button use requires recognition rather than recall memory, which may contribute to its greater usage than other revisitation mechanisms. Moreover, the Back button supports a commonly desirable movement in a hyperspace in a visually compact way [Cock+99, Gree+99].

Cockburn and Greenberg [Cock+99] have pointed to the Back button's inefficiency in retrieving distant pages as a potential limitation, and have further suggested that the stack-based rules it typically uses may not appropriately support user behavior. Additionally, the Back button's pull-down history list requires careful page title writing; many times, page titles are left ambiguous or receive little attention, which may be critical should users frequently engage in history list inspection. Importantly, both the mechanism rules and issues surrounding hyperlink presentation have been addressed as potentially critical for the mechanism's success in supporting user behavior.

3.7.1 Mental model

The Back button may seemingly be a rather simple mechanism, but navigators generally have naïve models of how it works [Cock+96, Cock+97, Gree+99]. Cockburn and Jones [Cock+96] demonstrated that even technically-oriented users were unaware of the mechanism's rules, generally believing it represented a simple linear list rather than a stack. The inability of users to recognize arguably simple mechanism rules (and indeed the failure to successfully investigate the behavior given the opportunity in the Cockburn and Jones [Cock+96] work) paints a grim picture for navigators' understanding of navigation mechanisms. The silver lining is that despite these generally naïve models, navigators make robust use of the Back button [Gree+99]. We will discuss mental models at greater length in Chapter 12.

3.8 Site map

Designers have a number of options for visually presenting the full (or, to varying extents, abbreviated) organization of a site, and most fall within either (i) content lists (such as with the "generic structure" of an index [Vora+94], or with a table of contents), or (ii) graphical representations. Both of these representations, when one considers the common Web vocabulary, fall under the term "site map," essentially an overview of the site structure.

While the study described in Chapter 15 focuses on textual navigational support, a wide range of graphical site structure representations have been proposed, designed, and thoroughly reviewed. Indeed, graphical maps can and do make the limitations of hierarchical representation less obvious; as McDonald and Stevenson [McDo+98] point out, "an index may not highlight the links between nodes whereas a map would make these relationships explicit." However, practical considerations such as page loading time often push graphic-heavy maps aside. Moreover, Feiner [Fein88] suggests another potential problem: "Although problems of … run-time may eventually be overcome, a fundamental difficulty remains. Although nodes may always be constrained not to overlap each other, the same is not true of arcs." The potential complexity of graphical hyperspace presentations remains a concern, since for even relatively simple connectivity in a Web space, such presentations can be overwhelming, and may less effectively support visual scanning for specific information [Pilg+99]. Finally, a graphical map tends to occupy a great deal more vertical and horizontal space than a textual overview, and thus regardless of its location on the screen, the map presents a real estate problem.

Contents lists do not avoid these problems altogether, of course. Hyperspaces can be and often are much too large for a complete site map to be useful, and so information designers must weigh the tradeoff between providing a broad view of the Web space showing a wide range of hierarchical relationships but running the risk of those relationships being lost in the shuffle, or providing a narrow view, ensuring that those relationships available will likely not be missed, but abandoning the opportunity to give a more extensive overview of the site structure.

Overview presentation involves a number of hyperlink presentation issues. First and foremost, an information designer must organize the content of the overview, namely the linkable text snippets, and decide who is related to whom, who is linked to whom, and so forth. These structural decisions are often made earlier in the Web development process than they perhaps should be [Lin+00]. Overview structuring and hyperlink ordering requires the information designer to take into account several potentially critical factors, including content relationships amongst pages [Jul+97], likely user interests given a particular location in the site [Niel00], and the balance of the tree, including how many nodes appear at each level, in each sibling group, and in what order [NormK91], amongst others.

The ordering of contents lists on the Web presents particular kinds of concerns due to the typically vertical nature of such lists. Consider the following small example:

- Study I
    - Experiment
    - Results
    - Discussion
    - Design Implications
- Study II
    - Experiment
    - Results
    - Discussion
    - Design Implications
- Study III
    - Experiment
    - Results
    - Discussion
    - Design Implications

“Study I,” “Study II,” and “Study III” link to sibling pages. Notice that their screen distance from one another (in this case, four text lines are between them) is dependent upon the number of children each has. The screen distance is much greater for siblings at higher levels in the site hierarchy, which in part may lead to the general advantage of collapsible overviews when the number of nodes becomes large [Chim+94]. Note also that each hyperlink has a “height” within the overview (the list position attribute, as discussed in Chapter 2), and that when fully expanded, top-level pages appear lower in the overview than all descendants of all siblings higher in the ordering. Even with a simple list such as the one above, other complexities arise regarding both behavior and design, some of which will be addressed in Chapter 6 and by the experiment described in Chapter 15.

4 Navigational Scheme

As Web design has matured, common ways of organizing and coordinating navigational mechanisms have surfaced in practice. These organizational schemes are primarily attempts to give coherence to the navigational support provided on a page or throughout a site or sub-site. Recall that design schemes, in a sense, may help give a Web space identity. A navigational scheme is a subset of the design scheme, helping define the support, structure, and connectivity of a space.

4.1 Typology

No serious attempts have yet been made to classify navigational schemes on the Web, likely because navigation design practice is still maturing, and even those schemes which appear common are not easily characterized. However, there are a few potential divisions and attributes that may prove useful in thinking about such a classification. As with navigational mechanisms, user behavior makes the distinction between the design scheme of a Web browser and that of any given site clearly a useful one, and may divide schemes according to their inter- or intra-site focus. Likewise, the extent to which a scheme is temporal or structural in nature will probably find its way into scheme classifications as they begin to show up.

The extent to which particular kinds of hypertexts are essentially linear has been analyzed [Bota+92], and may provide a useful line of classification for design schemes. In analyzing structural attributes of sites, Gillenson, Sherrell, and Chen [Gill+00] describe sites along a spectrum, from those providing great freedom in, or minimal control over, user movement, to those strictly controlling movement. Recognize, however, that a single site can have a number of design schemes applied to various sub-sites and pages. Moreover, an essential component of a navigational scheme is the particular form of presentation it takes on. In order to classify them, it will be necessary to account for the specific coordination and screen placement of the navigational mechanisms involved.

4.2 Display attributes

The display of navigational support, like all visualization, is subject to a few primary limitations, namely that only so much support can be supplied in the amount of available screen space, and that navigators can only process so much navigational information per unit time.

Several basic issues in the display of a navigational scheme deal with the orientation cues it provides. Thüring, Hannemann, and Haake [Thür+95] point to the necessity of indicating equivalencies between location units, for example between the current page and navigational aids in the scheme. A number of common practices do so via highlighting or deactivating hyperlinks to the current page, and may have important effects on behavior and models of the site.

Additionally, a navigational scheme can be largely isolated from or embedded in the page content. Haas and Grams [Haas+98], for example, point to the common use of the top or bottom of a page for top-level hyperlinks.

4.2.1 Mechanism presentation

Navigational scheme display is in part a behavioral grouping problem [Germ+00]: the site designer must organize the navigational mechanisms on screen in such a way that the navigator can make sense of their purpose and use them effectively. In another way, such display is in part a problem of task-action mapping [Howe+91]: the designer can attempt to define a mapping between a simple navigation task and the screen click behavior required to carry it out. These perspectives are in one sense prior to an analysis of the usage or sequence of mechanism scanning and clicking, and in another sense a product of it. Various logical displays of navigational mechanisms have been argued for in numerous practical guidelines (for example, Nielsen [Niel00] and Fleming [Flem98]) and can be more or less independent of expected visual scanning and navigation behavior. Moreover, consistency in interface display and associated cues may be the crucial factors in performance and usability of the display [Howe+91]. At another extreme, mechanism placement may be explicitly based upon frequencies and sequences of click-stream behaviors and usage [Sear93], for example, a placement promoting the efficiency of hub and spoke [Catl+95] behavior, which we will discuss in Chapter 10. Where mechanism design must account for the local search within particular modules or frames, scheme design concerns itself with the preceding activity: global visual search for the appropriate module [vanS+01].

Such global visual search behavior is sensitive to the specifics of the presentation. In warning of the dangers of excessive cross-referencing and associative links, numerous researchers, including Edwards and Hardman [Edwa+89] and Lynch and Horton [LyncP+97], demonstrate that the presentation of navigation aids has significant effects on a user's performance and perceptions.

The tension between display space limitations and the desire to provide navigational freedom (less control on Gillenson, Sherrell, and Chen's [Gill+00] spectrum) leads to a common problem in navigational scheme design: clutter. In providing additional freedom of movement, the navigator is instead impeded by an overload of options and potential occlusion of important or relevant information. Darken and Sibert [Dark+93] provide an insightful definition for thinking about the potential costs of each additional cross-referencing link used in a scheme: "A cluttered world is one in which the number of objects is so great that it obscures important landmarks or cues."

4.3 Contextual attributes

Navigational schemes are part of larger design schemes on Web pages, and exist in the context of other design schemes used within a Web space. Moreover, it is useful to consider the user's information need in thinking about design schemes, as this need can often drive the strategies she chooses during navigation.

4.3.1 Site schemes

The same sorts of approaches and concerns regarding the competition, coordination, and interaction between navigational mechanisms may well be applicable to design schemes. We know, for example, that being confronted with multiple organizational structures for a Web site can hinder the formation of a cognitive map [Edwa+89], but in some cases can improve performance [Conk87, Moha+92]. However, in general very little is known about scheme interaction. It is clear that navigators are continuously confronted with varying schemes and presentations, especially during inter-site navigation. Appropriately characterizing and measuring what navigators are confronted with and how such characterizations can help us understand or predict a navigator's cognition during the navigation session may be an important line of research for Web navigation in the future.

4.3.2 Task and strategy

The larger context of why a navigator is at a particular site in the first place, and the sorts of strategies she brings into the interaction, are important considerations for design schemes. It is not just that some design schemes are better suited to certain tasks, as is abundantly clear [Wrig+90, Rada+92, Roue92], or that designers must be concerned with a variety of potential user information needs and task types [Smit+97], or even that certain task types will drive a navigator's choice of navigation and support [Niel89, O'Day+93]. Rather, the navigational scheme itself might impact the sorts of strategies the user will employ, and potentially impact the structure of her tasks. In the area of information retrieval, for example, Campagnoni and Ehrlich [CampC+89] reported that the length of search result lists and interface design characteristics can lead a user towards different strategies.

4.4 Relationship

In Chapter 2, we saw that a defining feature of the hyperlink is the relationship it defines and identifies between a source and destination, and in Chapter 3, we similarly saw that navigational mechanisms can identify relationships between pages and between the hyperlinks themselves. Schemes are simply one level higher in navigational aggregation: they can identify relationships between pages, hyperlinks, and mechanisms.

We have gotten a taste of hyperlink and mechanism relationships in our discussions of competition and coordination. At any given moment in a navigation session, the pertinent relationships are often those between the user's current location and the space of available information. A navigational scheme identifies a subset of these relationships, in some cases providing mechanisms that consistently appear (for example, a mechanism for links to the top-level of the site), and in others providing specialized mechanisms for the particular current page.

4.4.1 Focus+Context

Recall from our discussion of site maps in Chapter 3 that Web spaces are often much too large for all pages to be connected to all other pages. The common techniques for addressing this navigational problem fall under the Focus+Context approach [Card+99], attempting to provide both relevant local context, and potentially orienting global context simultaneously without navigational overload. Fisheye approaches [Furn81, Furn86], for example, are based on the observation by Furnas that from any given location, the local context is seen with greater detail while more distal areas are seen with progressively less detail, ultimately reduced to major landmarks.

Combinations of local and global mechanisms within a navigational scheme are simplistic ("flat") versions of what fisheye views attempt to accomplish. Distal landmarks (top-level links) are often included along with more concentrated local context, in such a way that distal patches of a site are progressively less thoroughly represented given some current user location. Recognize that these flat, simple Focus+Context presentations avoid the potential problem of more sophisticated or extensive approaches, namely increased disorientation due to excessive distortion [Bart+95], which defeats the purpose of Focus+Context in Web navigation. Note the basic tradeoff, however: as Darken and Sibert [Dark+93] point out, in virtual environments a conceptually wide open space can easily lead to increased disorientation as well.

4.5 Hierarchical

In Chapter 2 we discussed the essential attributes of hyperlinks, one of which was the structural — often hierarchical — relationship between its source and destination pages. In Chapter 3 we saw the first level of hyperlink aggregation, in which often homogenous hyperlinks with respect to this source-destination relationship attribute are listed together as a single navigational mechanism. Common examples of these mechanisms were top-level, sibling, child, and breadcrumb lists. We now can see the next level of navigational aggregation, in which these common mechanisms are coordinated into a growing set of common practices in navigational scheme presentation, as in Figure 4.1:

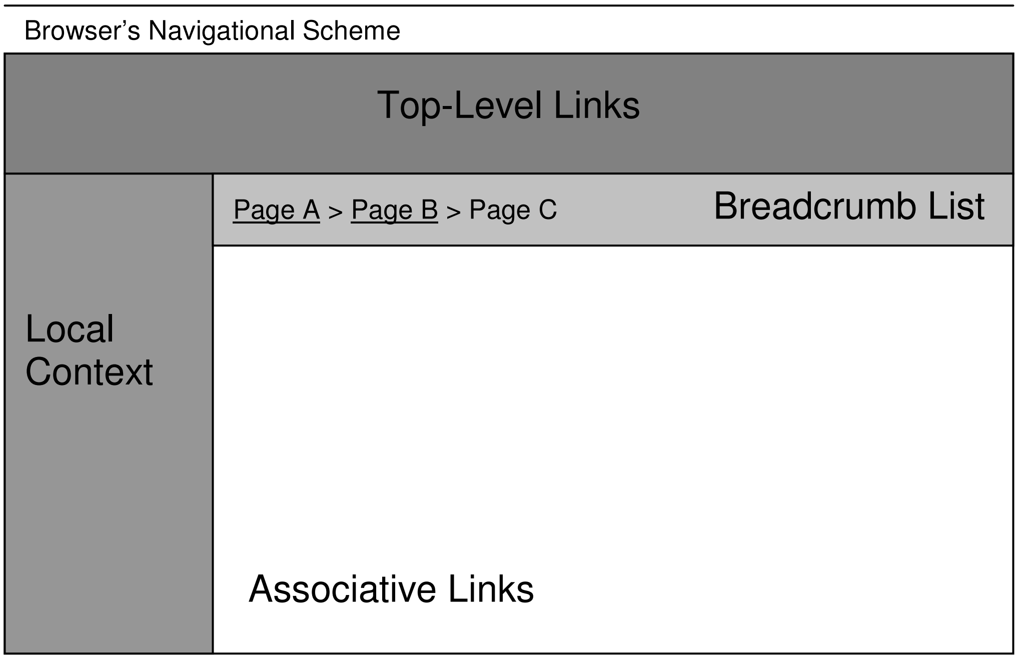

Figure 4.1: Common hierarchical navigational scheme. Set of top-level links horizontally presented across top of page, local context in the left column of page, breadcrumb list near the top of the content area, and associative links embedded within and/or separated from page content.

The practices depicted in Figure 4.1 arose as Web design matured. Presenting local context along the left column of the page with differentiated background color from that of the page content has been called "yellow fever" by Nielsen [Niel00], and its screen location, along with the common top-of-page placement for top-level links, have been found to best serve visual search performance [vanS+01]. There are other more fine-grained common practices, argued for by various authors with varying amounts of empirical support, for example the call by Nielsen [Niel99] for the convention of placing a hyperlink to a site's home page at the top left corner of every page within that site. We will return to and make use of these common practices when we discuss the study outlined in Chapter 15.

4.6 Homogeneity

To conclude our discussion of navigational schemes, we once again return to the notion of homogeneity of hyperlink attributes, in particular with respect to the interface changes that occur as a result of invoking hyperlinks within a given navigational scheme:

> *Scheme Homogeneity of Volatility:* the extent to which hyperlinks within a navigational scheme result in similarly volatile transitions.

Note here that given a set of homogenous mechanisms, the homogeneity of a navigational scheme simply relies upon the number of different mechanisms it coordinates, which is why schemes with very little distortion (or, put another way, limited mixture of global and local contextual navigation aids) are also the schemes with generally homogenous volatility.

5 Page

The page is the dominant metaphor in Web navigation, and its design, both for navigation and for content display, has possibly received more attention than any other aspect in the Web development process. The considerable attention has typically resulted in design guidelines, which, while potentially useful, often lack agreement and consistency [Ratn+96], may or may not be based on empirical work, and are notorious for typically being far too general for practical use.

We will approach the page, as we have approached the navigational constructs in the previous chapters, with a focus on definable attributes, empirical considerations, and, where there is little or no empirical evidence available, on theoretical considerations that may be empirically investigated.

5.1 Typology

Traditionally, pages have been classified by attributes of their navigational connectivity and content, with the two often explicitly separated for consideration. This split will be important throughout the remainder of our discussion, and has its basis in both design practice and navigational behavior.

Thüring, Haake, and Hannemann [Thür+91] suggest a structural classification of nodes based in part upon the amount of freedom of movement provided by the navigational support and the relationship between the source page and its linkable destinations. Structural nodes in their scheme are split into sequencing nodes, which are commonly used in highly controlled or linear hypertexts, and exploration nodes, which limit the imposed ordering. On the Web these nodes are often realized by site maps and guided tours.

The extent to which a page represents an aggregation of information available at other pages within or outside a site has typically been an important spectrum in page classification. Halasz [Hala87] first pointed to these aggregations as composite nodes in the NoteCards system, in which the "entire network would be 'included' in composite card, which the user could then use to refer to the network in its entirety." Such aggregation is the basis of site overviews and maps, and while they are obvious now, Halasz pointed out that they were not so obvious, and therefore missing, in the development of many early hypertext systems.

Botafogo and Shneiderman [Bota+91] point out that the aggregation of Halasz's composite nodes is only one of two critical abstraction concepts for node classification; the other is generalization, which "happens when a set of similar objects are regarded as a generic object," for example that a "hummingbird, a hawk, an eagle, etc., can be generalized to a single concept: that of 'bird'" [Bota+91]. Notice that the simple aggregation of Halasz's composite nodes assumes no specific attributes of the information space or the ways in which its pages are interconnected. Generalization, however, typically benefits most from particular types of semantic relationships, in particular those comprising IS-A or PART-OF hierarchies.

Botafogo and Shneiderman, in making their analyses, had the larger goal of formalizing many of the classifications for both connectivity and pages, and used the relative in-degree (extent to which a particular page is linked to from other pages in a space), and relative out-degree (extent to which a particular page points to other pages in a space). These notions potentially place index (high out-degree) and reference nodes (high in-degree) along a spectrum, but their analysis compares a node's in- and out-degrees to the means across a hyperspace.

Notice at this point a crucial split in the ways we can approach and think about Web spaces and classifications of their structures. Botafogo and Sheiderman take on the view of a hyperspace as a connected graph, classifying and identifying important properties according to connectivity, or who is linked to whom. This is fundamentally different from an analysis of the particular attributes of the hyperlinks on the page, such as screen location, which we discussed in Chapter 2. Recall from that chapter that embedded links are commonly thought of as links contextually included within a page's body text. As an example, suppose we have a page with a great deal of body text and several embedded hyperlinks. Such a page might have relatively high out-degree, but it may seem reasonable to think of such a page as clearly different from a site overview. The hyperlink presentation can give impressions to a navigator independent of the connectivity, and the page's connectivity is not always a salient feature. (Indeed, note that navigators are not made aware of a page's in-degree on the Web). We will see in Chapters 6 and 9 that there are problems with looking at a site solely in terms of its connectivity, and this problem may become more and more pronounced as navigational design becomes more and more structured and mature. Approaches that consider page design and layout have their relative advantages and disadvantages, but, crucially, appear better suited to discussions of user behavior and models of a site.

Haas and Grams [Haas+98] appear to have made the most effective effort to date in analyzing the characteristics of Web pages for a taxonomy. They offer the following characteristics as being important for such a taxonomy:

(i) The purpose or function of the page
(ii) Its intended audience
(iii) Its surface content or format (e.g., words, tables, sounds, tools, etc.)
(iv) The types of links it contains
(v) Its relationship to the pages to which it provides links (e.g., cover page, index, etc.)

Haas and Grams developed a page classification which included organizational pages, comparable to composite nodes. Organizational pages were sub-divided into tables of contents (listing the main sections of a site), site content pages (listing links within a site or sub-site, and may or may not be exhaustive), and index pages, generally with some sort of organizational strategy, such as alphabetical listings. Text pages were proposed as the broadest class, appearing to include most any information-centric page. Finally, the home page was distinguished as an introductory page to the site, and mostly "unique to the Web" [Haas+98].

5.1.1 Word count

The effort by Haas and Grams [Haas+98] was a movement towards classification schemes more explicitly considering design schemes, relationships amongst pages, and hyperlink attributes. In attempting to provide empirical support for page design guidelines, Ivory, Sinha, and Hearst [Ivor+01] showed how such considerations can lead to different conclusions about what makes for good design. In particular, they chose to group pages according to word count, with varying design metrics for the independent categories (for example, reporting that good low word count pages use less graphics, and good medium word count pages employ font variations between the header and body text). In their study, low word count pages had a mean count of 66.38 words, medium word count pages had a mean of 229.87, and high word count pages had a mean of 827.15 words. Whether there are behavioral or perceptual barriers (for example, the point at which content extends below the scroll line) beyond which additional content produces significant effects is unclear, and would make for interesting further research, in addition to better informing page classification schemes.

5.1.2 Centricity

Finally, we might characterize a page by the extent to which it is centrally located in the site, or the extent to which other pages might be easily reached from it [Rivl+94]. Such a classification deals primarily with the connectivity of the node, something we will discuss more in the next chapter. Rivlin, Botafogo, and Shneiderman [Rivl+94] propose metrics for a page's connectivity based on concepts defined by Harary [Hara59]. The status of a page is the extent to which it can reach other nodes in the space, while the contrastatus of a page is the extent to which other pages can reach it. Botafogo, Rivlin, and Shneiderman [Bota+92] further suggest that a page's landmark status might be determined by connectivity, for example by (i) its second-order connectedness, or the number of nodes it can reach in at most two steps, and (ii) its back second-order connectedness, or the number of nodes that reach it in at most two steps. These metrics correlate with the memorability of a node and are arguably good candidates for landmark status [Vald+88, Mukh+97], although they may not be predictive of behavioral landmarks [Modj+97].

5.2 Display attributes

The literature on Web page display is particularly extensive, coming in various forms, including practical design publications and empirical research reports. We will briefly consider a few attributes relevant to our discussion, and to the study described in Part III.

Web page display attributes fall under a variety of primary categories, including the extent to which a page displays primarily text, graphics, or a combination of the two [Thür+91, Card+01]. The screen density resulting from these interface objects is typically a primary concern as well [Dark+93]. On the flip side, we might consider the length of a page, or the extent to which it extends beyond the scroll line, and which interface objects end up occluded. Additionally, the particular layout and (often) modular structure of the navigational mechanisms and content area (or areas) helps define the page display in a useful way. Keyes and Krull [Keye+92] looked in particular at grid layout for pages, suggesting that navigators (i) visually distinguish chunks of information in the display (for our discussion, essentially scheme display), (ii) assess the relationships between the chunks (mechanism coordination), (iii) assign chunks to macro or micro status (for example, in global visual search), and (iv) distinguish further visual patterns and relationships within chunks (in local visual search, dependent on mechanism display).

The particular screen locations of the information chunks make up the crux of page display. While not yet standardized by any means, as we saw in Chapter 4, such display finds many of its roots in newspaper page design, the design of earlier computer interfaces, and in perceptual psychology.

5.2.1 Content and navigation

Notably, we make an explicit distinction between the content of a Web page, and its navigation support. There are two primary reasons for the distinction: (i) navigational support is typically presented isolated from the page content [Haas+98], and (ii) users scan, read, and react to link titles independent of the page content. Reason (i) is a matter of design practice, and alone would not be enough of a basis for the distinction; one could argue that such design practice might change over time. But reason (ii) has a behavioral basis, and is more fundamental; in a rapidly interactive environment like the Web, the potential navigation options are looked to independently of the information available at the current page.

5.2.2 Scanning and reading

A critical design aspect of a page is the extent to which it supports scanning behavior [Niel00], as users often spend little time on any given page [Cock+01]. Font variation to indicate relative headings and sub-headings, for example, is a widely made and validated [Ivor+01] design suggestion.

The importance of page design attributes supporting reading and scanning behavior will likely depend upon the extent to which a user is likely to read (and scroll) rather than scan. Neither of these issues has reached general agreement, as mentioned in Chapter 2. Regarding reading, the text width is an important attribute. In computer displays, lengths of 80 characters or more can be too long to be comfortable for reading [Lesk97].

In general, reading and scrolling behavior likely has much to do with the particular task the user is engaged in. Byrne, John, Wehrle, and Crow [Byrn+99b] reported frequent scrolling by navigators in a study using verbal protocols in the user's natural Web environment during generally undirected tasks. Other work has found scrolling requirements to often be a navigation barrier [vanW98].

Text length can have other implications as well. Diaper and Waelend [Diap+00] reported that users tend to base their at-a-glance page complexity ratings on the length of the page's text.

5.3 Contextual attributes

Web pages have content that is in relationship to the overall information scope of a site, and navigation options that, structurally and conceptually, place them within a subspace of that site. Often, this context cannot be seen by the user at a Web page; she is relatively unaware of where her location fits into a larger picture, a phenomenon known as the keyhole effect. Contextual information at least widens the keyhole and gives the page a more explicit relationship to other nodes. Eick and Wills [Eick+93] point to the importance of giving meaning to a node's relative position within the site. This can be tricky in developing a site architecture. Pages three levels deep in a hierarchy, for example, might be subjectively viewed differently than those at the second hierarchical level. Note that the relative position of an information chunk becomes more fuzzy, and additional complexities can arise, when a node appears in two locations within a site [Rivl+94].

Contextual considerations in interface design, including on the Web, are often dominated by the notion of consistency [Shne98]. A prevailing argument is that users will come to expect consistency in panel or screen real estate use to support their scanning strategies [Keye+92]. The extent to which a page's screen real estate usage and design scheme are similar to those of other pages in a space is a simple example of its contextual consistency. Homogeneity of content placement [Niel90a], for example, is especially common on the Web, driving or being driven by the user's tendency to look to the content when arriving at a destination page. Note an implication of such similarity for navigation strategies: in rapid backtracking, a user clicks the browser's Back button repeatedly until she recognizes the page she wishes to return to. The more similar the pages on the Back button's stack, the fewer distinct cues she has to quickly and easily recognize the desired page [Card+96, Cock+99]. Consistency, in general, is quite tricky and requires careful consideration from a number of perspectives.

5.3.1 Progression

Web pages are often either implicitly or explicitly part of a progression; recall from Chapter 3 that all hyperlink lists can imply some sort of ordered guidance. Progressions often determine the design attributes of several pages, in order to support the notion that navigation is restricted. Such context can often require more or less support of forward navigation [Thür+91], depending upon the extent to which a navigator will benefit from being relatively aware or unaware of what lies ahead in the designed progression.

5.3.2 Sub-site membership

Finally, a Web page's identity owes much to its membership within subspaces of a site. Sub-sites are generally information patches of tightly related topics. Their pages typically share a number of design attributes, such as branding, color coding, and similar navigation support, especially when the design scheme is locally focused.

5.4 Relationship

Progression and sub-site membership were examples of attributes speaking to the relationship between a Web page and other nodes within the same space. We can talk more generally of relationship spaces of a given page. The site space, for example, is simply all the pages contained in the site. All pages within a site presumably have some level of relatedness to all other pages; they must, after all, be related in some fashion, which is why they coexist. Pages can be thought of as having associative spaces, which are not explicitly shown to be related by hierarchical positions, and peripheral spaces, which are nodes relatively unrelated to the page in question. Some of the relationships between these spaces and common design practices on the Web are discussed in previous work [Dani01].

Note again that information topics can sometimes be presented in two separate hierarchical locations in a site, giving a single page, in some sense, two identities. For this reason, in some cases a page might be thought of as having multiple sets of any given type of relationship space.

5.4.1 Hierarchical

Familial spaces may borrow all of the traditionally used metaphors when speaking of hierarchies, including sibling, child, parent, grandchild, and grandparent spaces. Depending on the size and structure of the site, there can of course be many more such spaces. It is useful to consider the general proximity of particular spaces to a page, given general design practices, as discussed in Chapter 4. Siblings, children, and parents are commonly available within one click, the "immediate family" of the current node. Grandparent and aunt/uncle nodes can also be available in one click with many navigational schemes. However, the grandchild and cousin spaces of a page are rarely available within one click. Some of the potential implications of these common design practices are also discussed in previous work [Dani01].

5.4.2 Task and user

We will discuss task attributes in Chapter 7 and user attributes in Chapter 8. Note here that information topics explored during a navigation session must be thought of in the context of the user and her information need. The design attributes briefly reviewed in section 5.2 can be more or less appropriate given particular tasks and strategies. Some tasks, for example, will more readily call for quick scanning of a page and rapid movement in the site, while others will call for in-depth reading and extraction.

5.5 Residue

In Chapter 2, we saw that hyperlinks, as proximal cues to distal content, provide scent for information topics in the site; the hyperlink's text snippet can be more or less descriptive of the content it leads to. From another perspective, we can stand at a Web page looking out, and note that spread throughout the site is residue of the content available at that page. Furnas [Furn97] provides an analysis of the concept, stating that a page (or some target information chunk) has residue at a hyperlink if that link's attributes would lead the user to take the link in pursuit of the target. Furnas further points to the "daunting challenge" of navigation design: to structure a space in such a way that every node provides good (not misleading) residue for every other node.

6 Site

In this chapter we will discuss some of the primary attributes of a site related to our discussion as a whole. The study described in Chapter 15 will be concerned with relatively small sites, or sub-sites of larger ones.

6.1 Typology

Classifications of Web sites typically look to the site's purpose and content; news sites are distinguished from business sites, which are typically further distinguished from search engine sites. Classifying Web sites is a tricky business for a number of reasons, not the least of which being that as the Web grows, new purposes and new content domains tend to arise rather quickly. The study in Chapter 15 will be concerned with information-centric sites, where the purpose is simply to provide information on a (more or less) unified topic.

6.1.1 Purpose

A site's purpose, or in a sense what it provides for the user, is often a natural line of classification schemes. Shneiderman [Shne97] points out that we can often get at such a classification scheme (in a somewhat indirect way) by looking to the originator's identity, "which gives a quick indication of what the likely goals are and what contents to expect: corporations have products to sell, museums have archives to promote and government agencies have services to offer." Thus, originator's identity also clues us into the content domain of the site, another way of categorizing Web sites.

The purpose of a site, of course, may be viewed differently by its creators and by its users. In some cases, the creators' purpose might be as simple as providing pointers to other Web sites.

6.1.2 Domain

Site categorizations based on the content domain can perhaps be the most difficult to develop. The primary cause is simple: indexers and cataloguers of semantic content show great variability and disagree widely [Broo98]. The specific content domain, independent of the user's expertise with it (a topic discussed in Chapter 8) will not be our concern. But note here that content domains often lead to particular ways of structuring a site, often in such a way as to reflect the conceptual structure of the content; the domain, then, often drives the design scheme and structural organization. Some content domains, for example, lend themselves to IS-A or PART-OF structures. Moreover, the content domain alone can sometimes provide information about the purpose of the site, such as whether it is likely to be information-centric.

6.1.3 Size

Shneiderman [Shne97] provides an example taxonomy focusing on site size as a basis for classification. For example, he notes that a "personal bio" site might be between 1 and 10 pages, while a digital library site might be greater than 5 million pages in size. The typical sizes within particular genres of sites may change with time. Site classifications will likely benefit from more fundamental schemes, such as from our size discussion of navigational mechanisms in Chapter 3. As of yet, classifications, likely for pragmatic reasons, are centered around the era, or what is the common practice of the day.

6.2 Connectivity

A natural and rather common way of looking at a Web space is as a directed graph [Furu+90], in which Web pages are treated as nodes, and arcs connecting nodes are used to represent hyperlinks between the nodes. Brown [Brow89] has pointed out that representations of a real hyperspace as a directed graph can be mind boggling, and that, a priori, there is no hierarchical structure to a site. So information designers have a basic challenge: to effectively show the most important connections without stepping into overly complex territory, and to, when the site is meant to be hierarchical, give such an impression.

Botafogo and Shneiderman [Bota+91] point out that the number of nodes and links can give a rough idea of a site's complexity, but there are perhaps more informative metrics. Compactness indicates the extent to which a Web space is interconnected, while stratum indicates the extent to which the space is linear [Bota+92]. These notions have been referred to as "hyperchaos" and "drill and practice" [Marc+88], emphasizing the potential complexities of interconnectedness in a Web space. When sites are primarily linear (high stratum), they may be said to have a "spine" [Mars+89], from which detours are made possible. Finally, Gillenson, Sherrell, and Chen [Gill +00] note a common graph structure on the Web, which they refer to as a root return structure: all pages are linked to the foundation (or home page).

A number of other metrics are potentially useful, such as authoritative reference, when a set of nodes are all referenced by the same source [Chi+00], and the bi-directionality of hyperlink connections [Bota+91], although what we ultimately want to know is how these metrics play out in terms of user behavior and attitudes regarding a navigation experience.

6.2.1 Sub-sites

Connectivity can sometimes expose tightly bound subspaces of a site. Botafogo and Shneiderman [Bota+91] define a semantic cluster as a sub-graph of the site with a compactness level higher than that of the entire space, often making up a small Web locality [Piro+96]. As the authors acknowledge, a hypertext can be written in such a way that connectivity metrics will not tell the whole story. What is perceived as a sub-site, of course, may or may not match with connectivity attributes that a navigator is generally unaware of.

6.3 Consistency

Consistency in an interface begs at least two questions: (i) Consistent at what level? and (ii) Consistent with what? Both show the notion to be complicated; in our discussion, we will distinguish consistencies within a Web site from those between sites.

6.3.1 Internal and external

Structural and design attributes of a site can be consistent in a number of ways; the linguistic homogeneity we mentioned in Chapter 2 was an example of consistency within a site. Internal structural consistencies, for example, can occur when sub-sites are broken apart in similar ways. University Web spaces are good examples of this, as the department sub-sites often have at least a few similar hyperlinks, such as one for the department's course offerings and another for a list of the faculty. The structure of this thesis has a few obvious internal consistencies as well, for example the navigational construct chapters each have sections for taxonomies, display attributes, and contextual attributes.

When a site's content domain becomes common enough on the Web, external consistencies for that particular genre can begin to arise as well. University sites are again a good example, as a result of the organization of the content domain; universities tend to have the same sorts of departments. External consistencies can be more subtle, for example when across bookstore sites it is common for a hyperlink for an author's name to point to a list of books for sale by that author.

The recognition of these consistencies, as one might imagine, can be largely a matter of expertise, which we will discuss in Chapter 8. As a result of increased expertise with a topic, these internal consistencies within the expert domain may become more and more noticeable. As you read, you will hopefully begin to see more connections and internal consistencies in the overview of this thesis than you did when you first began reading it. If not, you may blame the author.

6.4 Homogeneity

We note again that, just as with mechanisms and schemes, an entire site can have more or less homogenous volatility transitions between its pages. The linguistic homogeneity mentioned in Chapter 2 was one example. We might also consider:

Structural Relationship Homogeneity of Volatility: the extent to which hyperlinks (across a site) with similar source-destination structural relationships result in similarly volatile transitions.

Directional Relationship Homogeneity of Volatility: the extent to which hyperlinks (across a site) with similar source-destination directional relationships result in similarly volatile transitions.

In many hierarchically organized sites, structure and direction can well determine the sorts of changes that will occur as result of following a hyperlink. For example, common navigational schemes often ensure that all parent and child nodes will differ in more or less the same way. As has been a theme throughout our discussion, however, that a site is connected or designed in a particular way does not imply that a visitor will notice, benefit from, or even behave as though she notices or benefits from, those structural and design attributes.

7 Task

The task is often what connects the navigational constructs we have discussed thus far, and the user, who we will discuss in the next chapter, and there are clearly complex interactions between task, user, and design scheme attributes [Niel89, Roue92]. Although the study described in Chapter 15 is concerned with a narrow strand of the task typology, namely individual fact-finding missions, it will be useful to briefly review the space of task types, and see where the concern of this thesis fits into a broader picture of activities on the Web.

7.1 Typology

Most Web navigation studies make a distinction between general purpose browsing, where the user consults sources with items of interest, and directed (or formal) search, where the goal is known [Catl+95]. The former often involves an undirected surfing behavior, in which the user sets out only to be confronted by something interesting [Byrn+99b]. Cove and Walsh [Cove+88] further suggest serendipitous browsing, where a user's movement is purely random, as a potential task type.

Browsing has been further decomposed by Marchionini [Marc95] into directed, semi-directed, and undirected browsing, with the classification focus being on the extent to which the user's activity is systematic and focused. Directed browsing in this taxonomy is quite similar to what many others call formal search, including tasks such as scanning for a known target, while semi-directed browsing is similar to what others refer to as subject-based exploration [Newf+98]. Notice that this classification includes consideration of the user's strategy and activity in carrying out the task, relying on some notion of "systematic" in assessing task type.

Shneiderman [Shne97], calling formal search "fact-finding," suggests two complex task types in addition to formal search and general purpose browsing: (i) extended fact-finding, where the user attempts to find more items like the object of an original fact-finding mission, and (ii) exploration of availability, where the user sets out to find all (or as much as possible) of the available information within a space on a particular topic, or whether a certain kind of information is available at all. Extended fact-finding is similar to the embedded task type [Sutc+98] where a particular task is motivated by some other external task.

Finally, Morrison, Pirolli, and Card [Morr+00] were interested in the sorts of Web tasks that had impact on users' actions and decisions, noting splits between finding facts, download items, and documents. They also distinguish between exploring (general browsing) and monitoring, where a user makes repeated visits to a site which updates its information. Note also a split between the task itself and the general notion of a user's purpose: Morrison, Pirolli, and Card recognized the use of the Web to find information, to compare and choose between alternatives, and to understand some topic. Moreover, notice that task complexity becomes a critical attribute. Some Web tasks require comparison of items across pages, or possibly across sites.

7.1.1 Formal search

Our primary concern in the study described in Chapter 15 will be with the formal search, or fact-finding, task type. Formal search may first be analyzed in terms of the number of distinct items the navigator is looking for simultaneously. Morrison, Pirolli, and Card [Morr+00] call multiple item searching a collect task, where the searcher is open to any answer. Note however that a user's focus may be in flux; when she believes she is close to some subset of her collection items, she might narrow what she is open to, and other times, switch focus amongst specific items [Jul+97, Neer+01]. In a laboratory setting, this can make data collection difficult to analyze since the goal of the user at any given moment is unknown (unless the user indicates her current cognitive state, but as we will mention in Chapter 10, this places the user in an unnatural setting). We will be more concerned with singular fact-finding missions, where "the granularity of the information is fine" [Choo+00].

Even singular formal search may be further decomposed. For example a navigator might look for a specific previously visited Web page (a go to task) or attempt to find a specific piece of information where the location is unknown (a locate task) [Byrn+99b]. Chimera and Shneiderman [Chim+94], in studying collapsible tables of contents, looked at tasks that involved finding information about the table of contents itself (not the site contents) and thus the object of inquiry appears to be an important way of distinguishing amongst formal search activity.

7.2 Strategy

From a standpoint in which observing user behavior on the Web is of primary concern, one of the most notable aspects of a task is the way in which it can drive the sorts of ways a user will choose to interact with a Web space. Wright [Wrig91] suggests, "task complexity rather than specific content is the major determinant of navigation strategy. The implications for evaluation are that without a detailed understanding of the tasks for which a hypertext will be used it can be very difficult to provide cognitively undemanding navigation options for readers." Recall from Chapter 5 that reading behavior may also be affected by the task type, with reading being prominent during undirected tasks [Byrn+99b].

7.3 Support

Tasks often lead to particular strategies in Web navigation, and Web navigation design can be more or less supportive of particular types of tasks. Supporting tasks is one of the greater challenges in information design, since a single Web space may often need to support a wide variety of user tasks [Smit+97], and different navigational schemes may be appropriate for each of them [Wrig+90]. Moreover, the information architect may rely in part upon considerations of how routine particular tasks are in designing for site usability [Frøk+00].

Some attention has been given to organizational structures and overviews and the circumstances under which they may improve performance. Shneiderman [Shne96] points out for example that some tasks, such as locating a previously visited page, or a node where the page title is known, are more clearly supported by indexical organizational structures. Overview structures have been shown to support performance in both formal search [McDo+98, DaniIP] and ill-defined tasks [Simp+89]. In general, structural context has been found to be supportive of directed search, and Park and Kim [Park+00a] claim that such context does not hinder general purpose browsing.

The presentation of the site structure is not the only concern; the underlying structure itself may be more or less supportive as well. Mohageg [Moha92] suggests that network structures, as opposed to more hierarchically structured spaces, may be detrimental to search performance, and Smith, Newman, and Parks [Smit+97] further claim that exploratory tasks are better supported by the less hierarchical network structures. The extent to which a site is linear, a topic we discussed in Chapter 6, is also a factor, with in-depth reading tasks likely being better supported by greater levels of linearity [vanN+99].

8 User

Characteristics of the user are critical considerations in the design of any interaction. When the general characteristics of a user population are known, a Web designer is at a great advantage. Under constrained circumstances, in fact, when the split between groups of Web site visitors is significant enough, presenting completely different navigational schemes may be advisable [Ersk+97]. Nielsen [Niel89], in analyzing several early hypertext studies, noted the prominence and importance of individual differences. Web navigators can differ in a variety of ways, including cognitive abilities and, for the analysis of their clicking behavior, as we will discuss in Chapter 10, the extent to which a given site will be explored [Luko+98]. Although such characteristics of the navigator are not the concern of this thesis, it will be useful to very briefly review important aspects of the user herself, if only to get a sense of the sorts of background factors she brings into the interaction with a Web space.

8.1 Typology

Web users do not easily fall into obvious classification schemes, primarily because, of all the major players in Web navigation, the human is by far the most complex. Shneiderman [Shne97], in discussing design issues for information abundant sites, suggests some of the obvious splits for consideration: "Gender, age, economic status, ethnic origin, educational background and language are primary audience attributes."

Probably the most discussed aspect of the user is her level of Web expertise. Commonly, navigators are classified along such lines as "expert to novice" or "expert to intermediate to novice."

8.2 Individual differences

Among the more interesting individual differences explored in the hypertext and Web navigation research literature are those relating to age and cognitive attributes such as spatial ability.

8.2.1 Age

Thus far, it is clear that younger Web navigators perform tasks more efficiently and encounter fewer navigation problems than older users. Younger users, for example, are faster information-seekers [Freu01] and are less affected by the hierarchical depth of a Web site [Zaph+00, Freu01]. Mental models, which we will discuss in Chapter 12, also appear to be affected. Neerincx, Lindenberg, and Pemberton [Neer+01] found elderly navigators to make more errors in a page categorization task, assessing the user's ability recall a site's structural organization.

8.2.2 Cognitive

It is clear that Web users bring a range of cognitive abilities and preferences to their interactions with a site, with spatial ability receiving considerable attention with respect to navigation behavior [Höök+96]. The propensity to become lost, for example, is related to one's spatial ability [Vice+88].

8.3 Expertise

A user's level of expertise, in a number of respects, is a critical consideration. Four types of expertise are especially relevant to Web navigation: general Web navigating expertise, content domain expertise, site domain expertise, and expertise with the currently navigated hyperspace due to previous exposure. Content and site domain expertise have subtle differences. Managing a bookstore, for example, might make one knowledgeable of how book stores in the real world are organized and run, but there is a difference between being knowledgeable of real-world bookstores, and being knowledgeable of how bookstore Web sites are typically structured and designed. Genres of sites tend to accumulate external consistencies that are not a priori known to even the expert of their content domain.

9 Transition

In the preceding chapters we reviewed the major players in the Web navigation game, looking at the navigational and organizational constructs that make up a hyperspace, the user, and the information needs she might bring into the interaction. The interaction itself — the process in which the user navigates through the Web space — is most often characterized by transitions from one page to the next, the focus of our discussion in this chapter.

In thinking about page-to-page transitions, it is useful to review the sorts of relationships that exist between the content and navigational schemes of the source and destination, some of which we already discussed:

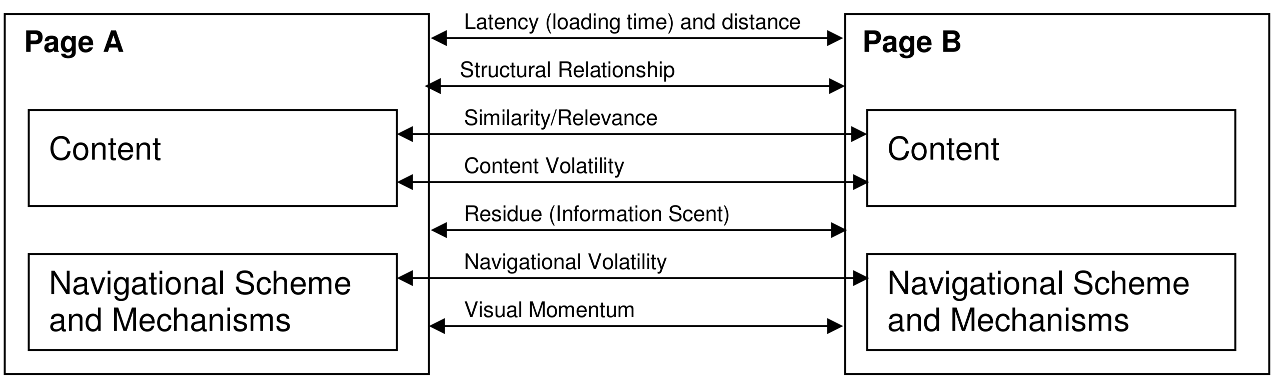

Figure 9.1: Page-to-page relationship. A transition from one page to another involves relationships between the content and navigational schemes of the source and destination.

Within hierarchical Web sites, we said in Chapter 5 that pages have structural relationship spaces, and that each page both provides scent to other information chunks and has residue spread throughout the rest of the site. Additionally, the content of two pages may be more or less similar, often indicated by or correlated with their hierarchical relationship. Moreover, the focus of our discussion and the study described in Chapter 15 is a characterization of the relationship between two pages based upon the changes that occur in the content and navigational options from the source to the destination pages of a transition. These will be discussed in Chapter 13. In this chapter we will look specifically at the latency and distance of a transition, its visual momentum, and the essential movement from one topic to another in page-to-page transitions.

9.1 Typology

There are two critical questions to be answered whenever a navigator moves from one Web page to another:

(i) What is the relationship between the source and destination pages of the transition?
(ii) Why did the user make the transition?

Given these two important questions, carving up the space of page-to-page transitions will require consideration of both the state of the Web space and the state of the user. The Web space gives us a classification scheme that should come as no surprise; for example we can think about transitions in terms of the hierarchical relationship between the source and destination pages, the hierarchical direction of the jump, as well as its distance, as we will discuss in section 9.3.

Transitions may also be classified by the navigational mechanism used, such as distinguishing Back button clicks from site map clicks. But there is a more interesting potential classification scheme, namely one that deals with the causal influences behind the transition, although these are less easily determined. Note that we can in some cases attempt to deduce the likelihood of particular user intentions by simpler classifications. For example, backtracking is generally a sign that the current navigational options carry low scent given the user's information need [Card+01]. However, a taxonomy will require an exhaustive set of user reasoning in making page-to-page transitions, which is not yet feasible given what little we know of moment-to-moment cognition in Web navigation.

9.1.1 Intention

The heart of a transition taxonomy, as alluded to above, will be a user's reasoning during the navigation process. At an intuitive level, we can think of a number of possible reasons for moving from one Web page to another, two of which might capture a wide range of user reasoning: (i) a proximal cue (hyperlink text snippet) appears promising given the user's current information need, so she invokes the hyperlink, or (ii) the current page and its navigation options do not appear promising given the user's current information need, so she attempts to move to another information patch. Particular needs at any given moment in a navigation session, such as the need for reorientation, might lead to certain kinds of transitions (as with backtracking, mentioned above), if the user recognizes the sorts of hyperlinks that will result in reorienting transitions. Whether navigators commonly gain the kind of predictive power needed to make movements for the purpose of fulfilling needs like reorientation is not known; other than backtracking and perhaps invoking persistent home page hyperlinks, it is not clear that Web users have sophisticated strategies for reorienting themselves.

9.2 Latency

Nielsen [Niel99] has proclaimed, "download speeds are the single most important design criterion on the Web." While it will not be a focus of the study described in Chapter 15, it is useful to note the prominence of latency in page-to-page transitions, and the flood of recommendations for short page loading times. These guidelines are often based on simple response time rules:

0.1 seconds: perceived as instantaneous by the user

1 second: activity is not interrupted, but the user notices the delay

10 seconds: the user will wish to perform other tasks while waiting for the system response

9.3 Distance

For every page-to-page transition, or succession of transitions, it is useful to think about the leap a navigator makes in terms of distance. Inter-site movements, for example, might be thought of as being of greater distances than intra-site movements.

The first, intuitive attribute to look at in thinking about distance is a site's connectivity, allowing us to define the distance from one page to another in terms of the shortest path between them. Utting and Yankelovich [Utti+89] call this the "minimal links placement strategy," and point out that other connectivity attributes, such as the number of hyperlinks two documents share, might reasonably be included in a distance measure.

Distance metrics, such as defined by Botafogo, Rivlin, and Sheiderman [Bota+92], are almost always defined at the site level, and therefore ignore page design, likely because accounting for page design is not easy in real world cases. Nonetheless, the concept is still crucial.

For starters, we might reasonably treat two pages with the same links as drastically different in the nature of their connectivity depending on the link attributes, such as location, for each of the links they contain, and the extent to which the links may be expected by the user at the time of the transition to the destination page.

9.3.1 Hyperlink attributes

Additionally, suppose we have the following two situations:

(iii) Page A contains a direct link to Page B. It is the only link appearing on Page A.

(iv) Page A contains a direct link to Page B. The link requires vertical scrolling to become visible, and is the i^{th} in a list of n links.

We may understandably wish to define the distance between pages A and B as being greater in situation (ii) than in situation (i). Three link attributes have been considered here: screen location, list position, and navigational competition. Navigational competition does impact the likelihood of a link being chosen, but it should not be taken for granted that screen location and list position do the same. Whether or not the scrolling requirement in situation (ii) will affect link choice depends on a user's willingness to scroll, and as we saw in Chapter 5 this is not a completely harmonious subject in the literature, with opinion and empirical evidence varying [Chim+94, Bach+97, vanW98, Byrn+99b, Niel00]. The effect of list position also lacks widespread agreement, if one considers the space of menu selection research. Recall, again from Chapter 5, that scanning of a menu may be random, top-to-bottom, or systematic in some other fashion [Card84, Hend89, NormK91, Horn+97, Aalt+98, Byrn+99a]. If top-to-bottom, the height in a vertical list should have an effect on link selection, and would therefore be a candidate factor in an analysis of page-to-page distance.

Another simple (and obvious) candidate factor for distance is the number of direct links between Pages A and B [Utti+89]; we might also consider the number of different ways a user can navigate from one to the other, which requires a consideration of all intermediate pages in all possible paths from Page A to Page B. Generally, distance metrics focus only on the shortest possible path and ignore all others.

Finally, one other contextual link attribute may need serious consideration: positional constancy [Teit+83], or the extent to which the user has already been exposed to pages in which a current navigational option appeared in the same screen location; this may be used in conjunction with the number of times the user has invoked the link in that particular screen location [Kapt93]. If indeed an important factor, we should accept the following result: suppose a home page link remains in the same screen location throughout a navigator's browsing session, and suppose the user visits Page A at t_0 and again at t_1. Then the user's distance to the home page will be shorter at time t_1 than at t_0 (assuming the links presented on Page A remain the same at the two different times, e.g. there is no within-page volatility via dynamic hyperlinks).

Given the potential importance of positional constancy, transitional volatility now has a privileged place in page-to-page distance metrics. At an intuitive level, pervasive positional constancy means low levels of volatility, and consequently user navigation experiences with high transitional volatility ensure that, in a candidate measure for the user's model of page-to-page distance, the navigator will view pages in the space as more spread apart. This will be a basis for hypotheses regarding perceived site size and global coherence proposed in Chapter 14.

9.4 Visual momentum

As mentioned in Chapter 1, visual momentum is a sort of antithesis to the kinds of changes a navigator might be rapidly confronted with while hyperlink following. The concept comes from techniques in effective film editing, and as Woods [Wood84] points out, it deals with those interface characteristics that support "rapid comprehension of data following the transition to a new display." Techniques for keeping an interface "flow" include, for example, keeping a set of top-level hyperlinks fixed at the top of the screen.

Simply stated, visual momentum is the extent to which interface characteristics from a previous state remain in relation to new interface characteristics after a user action has occurred, in some sense allowing the user to visually follow the change. Such visual momentum has been applied to hierarchical database navigation [Vice+88]. A classic example is an interface change in which a user opens a desktop folder, with the opened folder remaining in sight and its contents appearing indented below it. There is a common navigational mechanism on the web with the same characteristic: a (heterogeneous) sibling-child hybrid mechanism that provides a set of sibling links at one level in the site hierarchy, and another set of sibling links at one lower level, which are the children of the composite node [Hala87] for the section of the site the user is browsing. Notice that, particularly on the web, where latency is one of the primary areas of interest and concern [Niel95, Pitk98, Niel99], following the visual change may not be an easy task for navigators.

The focus here is the user's tendency to look to areas where important information will be available. Recall that navigators on the Web tend to look straight to the page content, particularly when the placement of page content is homogenous; the navigation options tend to be ignored initially, consulted only after scanning the content and deciding to flee from the current page. Thus, there is presumably a typical delay from the time of the transition to the time of being confronted with both persistent navigation options and new navigation options at the destination page.

9.5 Content

When a navigator follows a hyperlink, she is not just confronted with a potentially different visual display. She finds herself in the midst of different content. This new content may (i) be more or less related to the content of the source page, and (ii) be more or less well summarized or described by the proximal cue at the source page that led to it. We saw in Chapter 8 that different sorts of user expertise might contribute to the user's ability to predict these sorts of changes and potential surprises at the destination page.

PART II: MEASURING AND ASSESSING THE INTERACTION

10 Click-Stream Analysis

Whether in the lab or out in the wild, and whether for research or industrial purposes, the reporting of page accesses within and across Web sites has been a prominent way of conceptualizing and assessing navigation behaviors. Click-stream analyses typically involve (i) an ordered set of URLs (or some other code for documents accessed during a navigation session, and (ii) some kind of time stamp for each page visit. Research analyses typically distinguish between click-stream data for individual users. The first landmark click-stream study of Web navigation was conducted by Catledge and Pitkow [Catl+95], who analyzed 31,134 navigational acts taken from 3 weeks worth of data. Their study looked at global behavior, recording user events through a commercial browser.

When the data is collected remotely, it can often be skewed by network latencies [Ivor+01]. Moreover, what tasks users are performing during the navigation session is unknown. Lab work avoids these problems and allows for task and session attributes to be controlled, and moreover allow the researcher to control the complicated attributes of the stimuli (such as the site and its navigational scheme) we discussed in previous chapters. But we should note up front that even lab-based click-stream approaches do have their limitations. Ordered sets of page visits with time stamps are just that: ordered sets of page visits with time stamps. They do not provide a detailed look at the moment-to-moment cognition of the user [Card+01]. Using verbal protocols during the navigation session can help avoid this limitation, but presents a less natural navigation session for the study participant. Eye-tracking approaches provide more detailed information regarding user attention, but are not yet common in the research literature. Our discussion, and the study described in Chapter 15, will be most concerned with the sorts of ways we can characterize click-stream data, and what that data can potentially tell us about user cognition on the Web, especially under those circumstances in which the researcher can control task and session attributes, and enrich the click-stream data with more detailed information about the page-to-page transitions.

10.1 Paths

Perhaps the most widely accepted description of a common click-stream navigation behavior is that of the hub and spoke methods proposed by Catledge and Pitkow [Catl+95]. Such path behavior involves starting at a relative top-level node (the hub), digging along some path from that hub (creating a spoke), and then returning to the hub to create another spoke. In the “leave as you entered” strategy, users backtrack to the hub in order to create another spoke. This strategy is distinguished from a looping back behavior in which users return to a starting point by utilizing a navigation aid or browser feature. Note that with either strategy, “users rarely traverse more than two layers in the hypertext structure before returning to an entry point” [Catl+95].

Canter, Rivers, and Storrs [Cant+85] provide a set of path characterizations for movement through navigational structures. They propose the following measures in assessing a navigation session: (i) NV/NT, or the ratio of nodes visited to the total number of nodes in the space, (ii) NV/NS, or the ratio of the number of different nodes visited to the number of total node openings, (iii) pathiness, where navigation paths are such that no node is revisited, (iv) ringiness, where the path begins and ends with the same node, (v) loopiness, where the user navigates in a ring which contains no sub-rings, and (vi) spikiness, where the user returns to an origin node by tracing a path of previously visited nodes. Notice that spikiness characterizes one of the two behaviors ("leave as you entered") Catledge and Pitkow [Catl+95] described in proposing hub and spoke methods.

Trigg and Weiser [Trig+86] take a slightly different approach to thinking about path analysis, appearing to more explicitly incorporate potential user intention. In a movement along a train of thought, a navigator moves through a linear progression of topics. In a side trip, a navigator digresses from a train of thought, such as in following a hyperlink to the definition of an unfamiliar term within a hypertext before continuing. In a fork, a navigator decides amongst some number of possible trains of thought to take; in formal search, this decision presumably relies upon task attributes, or the current information need.

The sorts of fundamental path scenarios pointed out above, in conjunction with knowledge about the structural attributes of a hyperspace, allow for analyses of "patchiness" [Card+01] in Web navigation, or the extent to which navigators tend to stay within small, connected spaces for extended periods, as opposed to, for example, making frequent inter-site jumps. Tauscher and Greenberg [Taus+97a/b] reported the patchiness of Web navigation in their landmark click-stream study, noting that users tended to navigate in small clusters of related pages.

10.1.1 Routines

Paths and behaviors can be repeated, and there has been some interest in the sorts of routines navigators employ in the navigation process. Wright [Wrig91] points to a common reliance upon "superstitious mastery" or a tendency to make use of routine procedures, a behavior by no means limited to Web information-seeking. At a global level, Maglio and Barrett [Magl+97] report that navigators tend to repeat the same search patterns, such as using the same search engine.

These personal routines do not appear to extend to more fine-grained path behavior, however. Tauscher and Greenberg [Taus+97a], in applying a Longest Repeated Subsequences (LRS) algorithm to 6 weeks of navigation data from 23 users, reported that navigators generate only short sequences of repeated URL paths.

10.2 Visits and revisits

The metric associated with the ratio of the number of different nodes visited to the total number of node visits (NV/NS) proposed by Canter, Rivers, and Storrs [Cant+85] has received considerable attention in Web navigation research. Tauscher and Greenberg [Taus+97a/b] reported in 1997 that 58% of all page visits are revisits, based on data collected over 6 weeks from 23 users of a commercial Web browser, and point out that Web navigation activity therefore qualifies as a recurrent system, or one in which "users predominantly repeat activities they had invoked before, while still selecting new actions from the many that are possible" [Taus+97b]. Page revisitation was additionally examined in terms of frequency: 60% of pages were visited once, 19% were visited twice, 8% were visited 3 times, 4% were visited 4 times, and few were visited frequently. Tauscher and Greenberg further developed a recurrence distribution from the data, reporting a 39% chance that any given visited page would be a member of the set containing the previous 6 pages visited, showing that Web navigation involves a high degree of recency revisitation.

Cockburn and McKenzie [Cock+01] have since argued that the revisitation rate in global navigation is likely much higher, reporting that 81% of page visits are revisits, based upon data collected from 17 users over a 4-month period. They point out a number of potential contributing factors to the increased revisitation rate, including the evolution of Web navigation aids (as we discussed in Chapters 3 and 4), and the length of the navigation session being analyzed. The former is particularly interesting for our current discussion, as the effects of navigational schemes on revisitation rates have not been studied extensively. Site attributes such as size and connectivity might also contribute to these rates, and would make for interesting research topics.

10.3 Method of access

We briefly examined in Chapter 9 various ways of classifying the sorts of movements a navigator might make in a Web space, based primarily upon hyperlink and navigational mechanism attributes described in Chapters 2 and 3. When these attributes are controlled for or tracked, click-stream analyses can make use of them in reporting the sorts of ways Web pages are being accessed during a navigation session. For example, in early hypertext research, Hardman [Hard89] pointed to the frequent use of a particular mechanism, namely a home page hyperlink, presumably for reorientation purposes. By including (in addition to a list of page visits with time stamps) the particular mechanisms used and the hierarchical direction and familial relationship of a transition in a click-stream report, we can gain more fine-grained analyses of navigation behavior.

Catledge and Pitkow [Catl+95] provided the first steps in the direction of more detailed looks at methods of page access, reporting the rarity of Forward button use, and the predominance of hyperlinking and backtracking in Web navigation (combining for more than 90% of all navigational acts); other methods of access, such as URL typing, are simply rare. (Recall from Chapter 3 the cognitive explanations for this result.)

10.3.1 Backtracking

In their extensive study of Web behavior, Catledge and Pitkow [Catl+95] reported that 41% of all navigational acts were Back button clicks, a staggering result. Subsequent research by Tauscher and Greenberg [Taus+97a/b] reported the number at 30%. McKenzie and Cockburn [McKe+01], in reporting revisitation rates at about 81%, were unable with their data collection methodology to determine the method of page accesses. However, note that revisitation and Back button use are clearly intimately related, as Back button clicks are necessarily revisits.

10.4 Duration of visit

The very basics of click-stream data allow a researcher to look at the amount of time spent on a Web page; as we mentioned at the onset of this chapter, though, when data is collected remotely loading time cannot be controlled for, and may contribute to duration of visit analyses.

In general, Cockburn and McKenzie [Cock+01] have shown Web browsing to be a rapidly interactive activity, with fairly small amounts of time typically being spent at any given page [McKe+01]. Recall from Chapter 5 our discussion of scanning behaviors; rapid interaction is in contrast to other click-stream data suggesting online reading to be common. Of course, the particular tasks the user is engaged in are quite likely to contribute to the extent to which online reading is prominent, and we should note again that only in controlled settings can the user task be accounted for.

Typically attributes of the page are looked to as predictive of the amount of time a user will spend on a Web page, although the method of access (for example, if the page was accessed via the Back button) is also a promising approach [DaniIP].

Note here an effect of the moment-to-moment cognition limitation we mentioned at the onset of the chapter. We can attempt to infer user intention by time spent on particular Web pages, but such inferences are difficult to verify. For example, revisiting a page can be representative of a number of intentions, and we might like to split revisits according to intention: (i) short-term revisits, where the user returns to a page for the purpose of reorientation or in hub and spoke behavior, and (ii) long-term revisits, where the user returns to a page with the intention to performing extraction tasks, perhaps because she believes she may have missed important information during her previous visits. Looking for cutoff points in the durations of page revisits might give us a crude measure, but will nonetheless never evade the essential limitation of click-stream approaches.

10.5 Contextual effects

Click-stream data is often interesting in its own right. Some of the early landmark studies mentioned above provided insights into user behavior on the Web independent of particular concerns for task, user, and design scheme attributes. In some instances, however, the click-stream data was used to infer a sort of activity, for example in defining searchers as more likely than general purpose browsers to follow long, directed paths [Catl+95].

For our present discussion, what will make click-stream data interesting is its ability to tell us something about a user's performance and perceptions. This approach exposes a wide open area in Web navigation research. As McEneaney [McEn99] has noted, "for all the interest in navigational paths … there have been relatively few studies that have sought to examine the relationship between patterns of navigation and outcome measures."

Our discussion of the wealth of attributes in previous chapters was in part an attempt to lay a groundwork for the contextual effects one might look for in click-stream analyses. Aspects of the task affecting navigational mechanism choice, a user's level of expertise affecting her navigation behavior, and the amount of navigational support affecting the extent to which users backtrack are all examples of contextual effects in click-stream analysis. The space of potential relationships remains huge, to say the least, and many of the particulars will be discussed in relation to an empirical study in Chapter 17.

11 Disorientation

When Vannevar Bush [Bush45] speaks of being able to “reacquire the privilege of forgetting,” he is touching on the notions of disorientation and cognitive overhead. At any given moment in the navigation process, a user is prone to forget many things, not the least of which being where she is in the site and how she got there. The purpose of navigation design, in part, is to allow her to forget without consequence — to be able to easily pick what she needs out of the information currently available to her.

Disorientation is a common phenomenon on the Web. Disoriented navigators have trouble gaining an overview of the information structure before them [McDo+98], and are often unable to find previously visited pages and desired information in a hyperspace [Park+00b]. Recall the three fundamental questions navigation design answers for the Web user [Niel00]:

(i) Where am I?
(ii) Where have I been?
(iii) Where can I go?

We can think of disorientation on the Web as a lack of understanding regarding (i) and (iii) above, and arguably regarding (ii) as well, although that sort of understanding will be discussed more in terms of a user’s mental model of her interaction with a site, in the next chapter.

Recognize that if we ask a Web navigator to describe where she is within a site, she could have two fundamental types of responses. She might refer to her model of the organizational structure of the site, saying something like, “I am at the home page” or “I am at the top level of the site, which is near the home page.” Or, she might refer to the information structure of the site, saying something like, “I am in the part of the site that talks about Italian restaurants in Palo Alto.” Both types of responses would be consistent with the common location-based metaphors on the Web [Matl+96].

As we saw in Chapter 9, transitions have both structural and informational characterizations. The user’s current location can be thought of in the same way; navigators move amongst information topics, which are near to or far from other information topics. But recall a somewhat surprising finding from Chapter 8: users appear to become lost independent of their content domain expertise [Elm+85]. This striking finding may give us reason to think about disorientation in terms of site structure, connectivity, and design, without as strong a concern for the specific information topics being navigated, but the relative contribution of web site structure and content to disorientation is an open area for research.

11.1 Lost in hyperspace

One of the signature phenomena regarding disorientation on the Web is the feeling of lostness many navigators are prone to. When navigational difficulties arise and a user needs to reorient herself within the space, she is often said to be “lost in hyperspace” [Edwa+89]. Lostness within virtual information spaces has been a fundamental problem for many years [Robe+81], and can often lead a navigator to temporarily abandon her primary information-seeking task in favor of reorientation strategies. The specific causes of lostness on the Web are relatively open for investigation; Otter and Johnson [Otte+00] found that users most commonly report ambiguous link titles as the cause of lostness. Notice the potential relationship between predictability and lostness: poorly labeled links most notably give rise to difficulty in predicting the content at the destination page. Notice that, in potential contrast to the structural causes of disorientation mentioned in the previous section, this reported cause of lostness by the users themselves does in fact have much to do with the information topics in the site, as it involves the inability to interpret hyperlink titles within a content domain. The extent to which structural and content domain factors interact to bring about disorientation problems on the Web also remains an open concern for research.

11.2 Cognitive overhead

Without the help of the information designer, a navigator must simultaneously focus on remaining oriented in an information space and on achieving her goals. Note here that keeping oneself oriented in a space is not a primary task. We saw in Chapter 7 the sorts of goals users have in navigating on the Web; avoiding disorientation was not one of them. The user who wishes to locate specific information in a Web site and additionally needs to focus on staying oriented in the space is experiencing cognitive overhead. Conklin [Conk87] used the term to refer to the additional effort and concentration necessary to maintain several tasks or trails at one time. One of the primary purposes of the various navigational constructs discussed in Chapters 3 and 4 is to significantly reduce cognitive overhead, and allow users to focus on their goals, returning to orientation information only when necessary.

Maintaining many potential paths within a hyperspace has been a primary source of cognitive overhead on the Web. Backtracking in a space, for example, can often be difficult precisely because the context of the current page is often quickly forgotten [Dufr+97]. Multiple window coordination presents a good example of such difficulty as well. Nararro-Prieto, Scaife, and Rogers [Nava+99] claim memory difficulties are pronounced when the user attempts to remember the context of more than three Web browser windows. The problem of overhead and ease of forgetting context can have disastrous effects in the Web navigation process, and is essentially a problem of limited mental resources [Albe97]. Often, entire sub-sites a user had planned to explore are ultimately neglected due to forgotten intentions [Newf+98]. Wright [Wrig91] points out that overhead problems can be aggravated by departures from linearity in hypertext navigation.

11.3 Metrics

Because we cannot "dig out" the proposed feelings of lostness described above, a few research efforts have been aimed at quantifying either the phenomena itself or possible causes of it. It is clear that thus far the approach of choice lies in analyzing the extent to which a navigator deviates from an optimal path to some piece of desired information. (Notice that this is a click-stream approach, making use of paths of node visits.) Mohageg [Moha92] used a measure of the deviation from an optimal information-seeking path to assess efficiency, and Smith [Smit96] later argued that we can think about disorientation in terms of these deviations. Otter and Johnson [Otte+00] followed suit, extending Smith's metric to account for hyperlink attributes that might contribute more or less strongly to disorientation, pointing out that "links involve much more complicated theoretical and design issues than at first appear." We discussed many of these complicated issues back in Chapter 2, noting for example that purely associative links may reasonably be viewed as less predictable than structural hyperlinks.

Recall our discussion of page-to-page distance in Chapter 9. We said there that hyperlink attributes such as screen location and list position could reasonably be incorporated into an assessment of distance; the number of necessary clicks between pages was viewed as likely too simplistic. Indeed, many ways of characterizing hyperlinks and the transitions that follow from invoking them are going to be crucial in thinking about both the distance between two pages (an important concept in the metrics proposed by Smith [Smit96] and Otter and Johnson [Otte+00]) and disorientation. Take for example a constantly full overview of a site's contents. With such a prosthesis, the optimal path to any information chunk from any other will involve one click. In cases such as this (and indeed cases close to it, where the compactness [Bota+92] is high), we will clearly need to look more closely at design aspects of the interface to understand the lost in hyperspace phenomenon and its contributing factors.

11.4 Prostheses

Hardman [Hard89] very early recognized that a hyperspace's home page was often a special, reorienting area. Context, or at least the ability to quickly regain context, is not surprisingly the most useful prosthesis for disoriented navigators; while the user is always at a single page, looking through the narrow keyhole, sensemaking requires the recognition of many [Card+96], only supported by navigational context.

Maps are a common prosthesis, but on the Web are generally presented isolated as an individual page, and such occlusion can greatly reduce usage [Wrig91]. Overviews in general enforce location awareness [Chim+94] and reduce the cognitive overhead required for accessing information within a site [Thür+95].

12 Mental Models

It is widely believed that navigators form mental models of the Web sites they visit, and further that they form models which are often spatial in nature. A mental model is an internal, symbolic representation of some part of the external world [John83, John91], and when applied to computer and information systems, Halasz and Moran [Hala+83] take such a model to be "a cognitive representation of the system's internal mechanics," which includes "its component parts and their behaviors." With respect to intra-site Web navigation, the external world is the Web site, the component parts, for the most part, are the navigational constructs we have discussed in previous chapters, and their behaviors include the sorts of changes they produce when invoked, per our discussion in Chapter 9.

12.1 Cognitive maps

The extent to which internal representations of information spaces might be spatial has received a good deal of attention. Such spatial representations typically come in the form of what Tolman [Tolm48] referred to as a cognitive map, and include route knowledge about particular paths through a space, landmark knowledge, and, when more fully developed, survey knowledge about the space's general landscape. Maglio and Barrett [Magl+97] have argued that user behavior suggests Web navigators conceive of sites in the same sorts of ways they conceive of physical spaces. Farris, Jones, and Elgin [Farr+01] have questioned whether users form spatial representations of the Web, which is in a sense inherently non-spatial. Our concern will not be with an ongoing cognitive map debate, but rather with the ways we might assess a user's model of a Web space and her own interaction with it.

12.1.1 Landmarks

Landmarks (or anchor points), or salient features of a space, are critical components of a spatial mental representation; moreover, Canter, Rivers, and Storrs [Cant+85] pointed to the importance of their identification in complex navigation spaces. Even if a navigator's representation of a site is not spatial, evidence suggests there are at least critical pages, both behaviorally and cognitively, in the navigational process that serve as anchor points [Shum90, Magl+97]. These anchor points may serve as reorienting destinations, such as a site's home page. Maglio and Barrett [Magl+97], moreover, have argued that navigators rely upon these anchor points to structure their memories of a Web space. It turns out that these memorable landmarks are often pages with high back second-order connectedness [Vald+88]. However, there is question as to whether such structural landmarks actually correlate with behavioral landmarks; Modjeska and Marsh [Modj+97] claim structural properties of landmarks are poor predictors of the kinds of behaviors one would expect of landmark usage. We should note here that many proposed structural characteristics of nodes in a Web space, such as back second-order connectedness, may be difficult or even more or less impossible for a user to notice, so the claim need not be that the navigators pick up on the connectivity of a Web page as a salient feature and choose to use it as a landmark page. Rather, connectivity attributes may over time tend to lend themselves to use as landmarks.

12.2 Hierarchies

We noted at the onset of our discussion in Chapter 1 that hierarchies would play a prominent role throughout this thesis. Put simply, hierarchies often tend to match the way we think about the world. This includes our thinking about navigational spaces [Stev+78, Chas83]. Modjeska and Marsh [Modj+97] have shown mental representations of Web sites to often be strongly hierarchical, as drawings of site organizations tend to show biases towards tree-like structures. In fact users' internal models tend to be hierarchical even when the site itself is not [Modj+97]. Dillon, McKnight, and Richardson [Dill+90] have claimed that such mental models enable navigators to more easily form cognitive maps. Finally, we note a relation to Web page memorability: Charney [Char87] has shown that nodes near the top of hierarchical hypertext structures are more likely to be remembered.

12.3 Site models

Mental models on the Web are primarily *about* sites. We can think of a navigator as forming a model of the sorts of attributes we discussed in Chapter 6, such as the site's connectivity. We will be concerned in our current discussion not only with models of the site in and of itself, but in a user's way of thinking about her own exploration of the site. A navigator can benefit not only from knowledge of which pages are connected to which, but from knowledge of the extent to which she has followed hyperlinks and explored sub-sites. Recall from Chapter 2 one of the primary navigation problems in hyperspace: the inability to determine or in a sense "see" entire paths of pages beyond the adjacent destination nodes of hyperlinks. In this thesis we will consider more accurate models of one's own exploration of a space as potentially alleviating that problem.

12.3.1 Size

As we saw in Chapter 6, a signature site attribute is its size. Modjeska and Marsh [Modj+97] have made empirical and theoretical claims about site size perception based upon exploratory research of navigation behavior, finding that strongly hierarchical sites tend to be perceived as smaller than non-hierarchical sites. They further speculate that greater numbers of navigational options in weakly hierarchical sites give a sense of extent and range, while the converse would be true for strongly hierarchical sites.

We will distinguish a user's perception of a site's size from her perception of a site's complexity, although the two are likely somewhat related. Larson and Czerwinski

[Lars+98] point out that a site's perceived level of complexity correlates with its hierarchical depth.

12.3.2 Global coherence

Internal representations of a site can be in relation to both the site's structural organization and its content, although these representations can and perhaps should often overlap [McKn+91]; that is, it is often argued that the structure and connectivity of a site ought to reflect its semantic content. One of the critical obstacles in successful information-seeking on the Web lies in achieving comprehension of a Web space's content and understanding the interconnections of that content, which is itself arguably a mental model construction problem. As Thüring, Hannemann, and Haake [Thür+95] point out, "in cognitive science, comprehension is often characterized as the construction of a mental model that represents the objects and semantic relations described in a text."

Comprehending a document is related to the notion of understanding the connections between its various information chunks. Understanding small-scale connections between sentences within a Web page demonstrates local coherence of the Page's content, while understanding large-scale connections between many nodes within the space demonstrates global coherence [Thür+91].

Global coherence often allows for forward directed inferences about a hyperspace, for example understanding what information is likely to lie ahead without having already seen it [Thür+91]. Forward directed inference also may come out of the extent to which a Web site has high levels of internal consistency, as we discussed in Chapter 6. In a sense what global coherence might improve is a user's level of predictive power in hyperlink following.

12.3.3 Exploration

Models of one's own exploration of a site are necessary for successful navigation. Web users often need a sense of the paths they have already thoroughly explored in order to make wise information-seeking decisions. Mental models of such exploration are important combatants to disorientation in hyperspace [Cock+96]. An understanding of exploration, we might expect, could support backward navigation, the ability to return to previously visited Web pages, or "finding information you have previously encountered" [Thür+91]. Recalling from Chapter 10 the extraordinary level of page revisitation that typically occurs in Web navigation [Taus+97, Cock+01, McKe+01], we can see how such support could be crucial. However, note that the extent to which accurate exploration models would be needed to support backward navigation is unclear, since most revisited pages have been seen within the last 6 page visits, per Tauscher and Greenberg's [Taus+97a/b] distribution. So, either the Back button may be all that is needed for successful backward navigation, or users return to recently visited pages only because browser history lists need serious design improvements, and the Back button by its nature supports recency revisitation — another chicken and egg problem.

12.3.4 Design

Users not only build up representations of the site content, but additionally form models of the site's "component parts and their behaviors" [Hala+83], as we mentioned at the beginning of this chapter. In other words, they form mental models of the navigational mechanisms and schemes. We talked about the naïve mental models users have of the Back button [Cock+96, Cock+97, Gree+99] back in Chapter 3. Successful navigation may in part involve the construction of coherent and accurate models of what sorts of ways the user's location will change when she invokes certain hyperlinks, or uses certain navigational mechanisms. The idea is for the navigator to come to understand a mechanism as a coherent entity, with a purpose all its own, as we said in Chapter 3. The same may be said of schemes. Recall that multiple organizational structures hinder navigation and can make a hypertext less usable [Edwa+89]. The theorized cause of this hindrance has been that multiple structures (with separate organizational displays) can hinder the formation of a cognitive map of the hyperspace [Edwa+89]. But if users also attempt to form mental models of a site's navigational scheme, a plausible explanation would also be that many different organization displays hinder the formation of such a model; the user is forced to work harder to understand the behaviors of the various mechanisms available to her.

12.4 Measuring and assessing

As with disorientation, we cannot peek inside the user's head to see her mental models of a Web site. Instead, a few tactics have been employed to gage the user's view of the site structure and some of the other types of models we have discussed in this chapter. Which navigational behaviors correlate with such models has received very little interest, although McEneaney [McEn00] has found that the total and distinct number of pages a user visits during a navigation session are not predictors of comprehension.

To investigate user understanding of site structure, van Nimwegen, Pouw, and van Oostendorp [vanN+99] had participants do a walkthrough of the site, attempting to visit all pages in a systematic manner. Their level of success was measured by their ability to do so, and by the number of superfluous pages they needed to visit in completing the task (given some optimal way of walking through all pages in the site). The assumption was that structural understanding would allow a navigator to better survey the site systematically. McDonald and Stevenson [McDo+96] have looked more to memory for a sense of mental model formation, by having participants attempt to locate nodes within the site after the navigation session. These sorts of approaches are markedly distinct from what many researchers do to investigate mental models of site structures: ask participants to draw the site.

12.4.1 Cognitive cartography

Cognitive cartography refers to a research practice in which experimenters ask participants to sketch a map of a Web site, generally including boxes for Web pages and lines for hyperlink connections. Although the analysis in this practice is often qualitative, Otter and Johnson [Otte+00] have used the technique to produce mental model metrics, looking at the accuracy and completeness of such drawings. Farris, Jones, and Elgin [Farr+01] have argued, however, that such sketches generally reflect the user's understanding of the semantic structure of a site, not its organizational structure.

12.5 Prostheses

Recall from the previous chapter that navigational prostheses can benefit users substantially, often reducing disorientation and increasing usability. These spoils, however, appear to come with a price. For all their good, prostheses introduce a potentially disastrous problem: they tend to make us lazy. Early in the short life of Web navigation research, it looked as though navigational prostheses could enhance mental model formation, helping reveal the structure of the site and demonstrating the conceptual structure of the content domain [Rada+92, McDo+98]. However, site overviews and other structure-revealing prostheses may not improve structural mental models [Dias+97]. In fact, Farris, Jones, and Elgin [Farr+01] claim that they can hinder such formation, as they may cause the user not to need to work for structural understanding.

13 Volatility

In Chapter 1 we briefly looked at the ways a site could be volatile, splitting the changes that a user might be confronted with into those which occurred within pages (for example, the home page of a news site updating its content daily) in Newell's [Newe90] time scale of human action, and those which occurred across pages as a result of page-to-page transitions in hyperlink following. The latter we referred to as transitional volatility, which we will now shift our attention to in this chapter. We will pay special attention to how such volatility might contribute to disorientation on the Web, or more precisely to that subset of disorientation that results solely from page-to-page transitions.

13.1 Transitional disorientation

Changes in content and navigation options from one page to another contribute to the whole of transitional volatility. What in part makes the concept important to a user's experience on the Web is its potential contribution to the lost in hyperspace problem. We will model transitional disorientation as that subset of disorientation in hyperspace resulting solely from the act of hyperlink following. The concept is modeled according to proposed factors contributing to need and potential for reorientation into the destination page. These factors are shown in Figure 13.1.

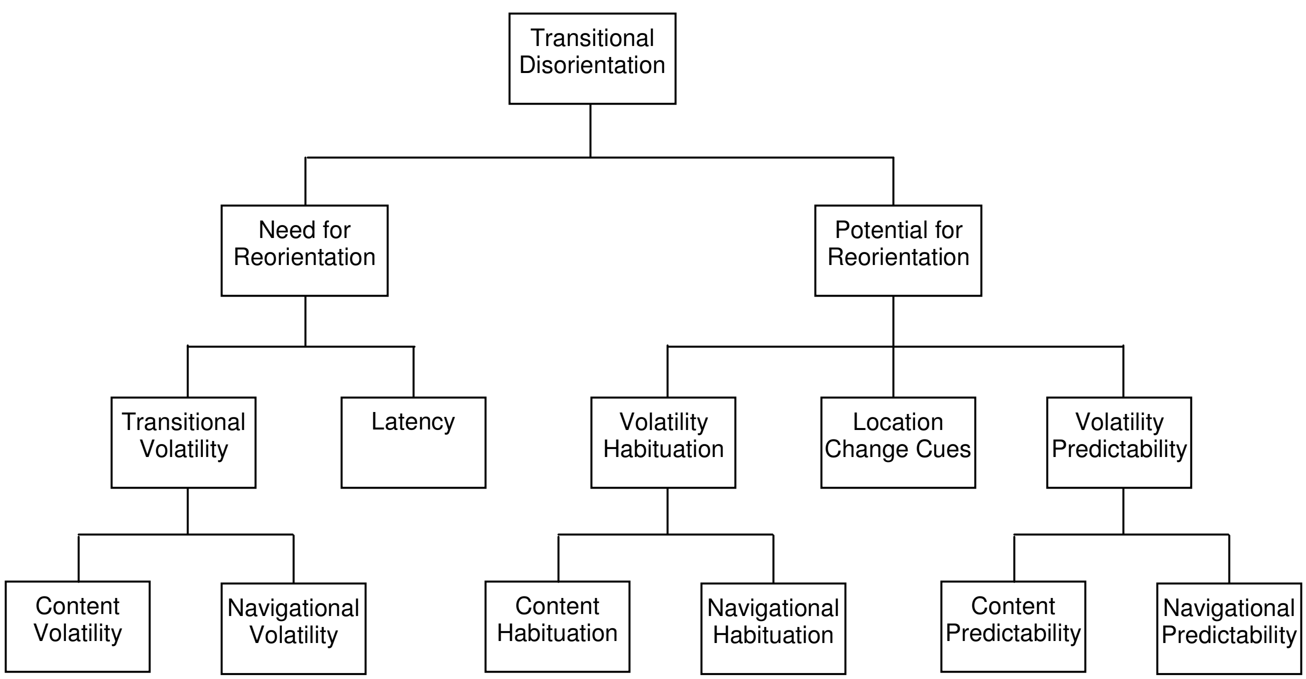

Figure 13.1: Transitional disorientation model. Descendants in this hierarchy are proposed as factors contributing to their ancestors. The content and navigational changes that occur as a result of a transition from a source to a destination Web page is modeled as a contributor to the lost in hyperspace problem, of which transitional disorientation is a subset.

Transitional disorientation results from two contributors, generally in a sort of tug of war. Need for reorientation refers to the extent to which the navigator requires reorientation at the destination page, while potential for reorientation refers to the extent to which the navigator is in fact equipped to reorient herself.

As a contributor to the user's need for reorientation, latency is simply the loading time of the destination page, or the amount of time between the invocation of the hyperlink to the moment of the destination page's rendering. (On the Web, page rendering is generally complicated, since some interface objects will be rendered while others are still being downloaded, but this issue will be accounted for in the study described in Chapter 15.) As a contributor to the user's potential for reorientation, location change cues are simply the differences between the source and destination page orientation cues. Both of these factors are fascinating in their own right, and could easily be the devotee of an entire thesis, but our concern here will be with the navigational halves of each of the three remaining sub-trees in our hierarchy from Figure 13.1. In particular, we will discuss:

(v) The interface differences between the source and destination pages' navigational support (navigational volatility)
(vi) The extent to which a user is habituated within a navigation patch with its own restricted navigational support (navigational habituation)
(vii) The extent to which a user is able to predict navigational support changes at the destination page before a page-to-page transition (navigational predictability)

13.2 Transitional volatility

Recall the simple notions of visual turbulence, visual momentum, and basic site volatility that we encountered at the onset of our discussion. Transitional volatility was said to relate strongly to those notions, with the extent to which a Web interface changes during the navigation process being potentially unique and fascinating for a number of reasons. In Chapter 3, we looked at a navigational mechanism with highly disparate destinations, some likely to look quite different from the source page (such as with a hyperlink to a different site), and others likely to look relatively similar to the source page (such as with a hyperlink to a page within the same sub-site). That is, we saw that the changes could be more or less drastic.

The interface changes that occur as a result of a page-to-page transition can be separated into those related to the content of the source and destination pages, and those related to the navigational scheme of the two pages. Differences in content are two-fold: they are both semantic and visual. In shifting focus from one information chunk to another, the user is confronted with new content at the destination that may be more or less related to that of the source page. If we assume the site is organized in such a way as to place related content close together, the extent to which the content at the two pages are different can be expected to correlate with the distance of movement, per Chapter 9; recognize that a user's characteristics, such as content domain expertise, are likely to play a role in her ability to recognize these differences. The visual content changes are the differences primarily between page display attributes of the two pages, many of which were discussed in Chapter 5. For example, the source and destination pages might have different color coding, graphics, fonts, and other text display features.

What we will focus more intently on in this thesis are the changes in navigational support that occur as a result of a page-to-page transition, namely navigational volatility. Navigational volatility may be decomposed into its hyperlink and mechanism components. That is, a movement in hyperspace potentially involves both changes at the micro level where the source and destination pages differ in some number of hyperlinks, and at a macro level where entire mechanisms shift and change during the user's navigation experience.

A simple measure of hyperlink volatility counts the number of hyperlinks appearing on the destination page that did not appear in the same screen location on the source page. Of course, this simple metric might be extended to include all possible source-destination screen real estate relationships:

(i) Hyperlinks on the destination in screen space occupied by the same hyperlink on the source page
(ii) Hyperlinks on the destination in screen space occupied by a different hyperlink on the source page
(iii) Hyperlinks on the destination in screen space unoccupied on the source page
(iv) Hyperlinks on the source in screen space unoccupied on the destination page

In assessing need for reorientation, we might favor measures referring to navigation options at the destination (such as our initial simple measure), where the user is required to reorient herself. Note also that (i) above is in fact a metric for visual momentum (at a micro level perceptually).

Looking at a transition in a vacuum and counting hyperlinks in this way says nothing, however, of the user's navigation history, which will be our focus in the next three sections.

13.3 Volatility habituation

If we wish to characterize a user's navigation session according to the extent to which source and destination pages contain different content and navigational support, we can note as part of that characterization those stretches of page-to-page transitions during which little or no transitional volatility occurs. In such stretches, a navigator might become habituated. That is, she might come to expect a lack (or low level) of transitional volatility as a result of a page-to-page movement based upon a recent lack (or low level) of transitional volatility. Again, her level of habituation deals with both content and navigation.

Habituation in a Web space might rely on a notion of recency in navigation. There are two types of recency in Web navigation: (i) the last *n* pages visited, and (ii) the last *n* seconds of interaction with a space. So, for example, we might measure a user's level of habituation at any given moment in a navigation session as the mean transitional volatility over the last *n* transitions. But how would we pick the right *n*? Tauscher and Greenberg [Taus+97a/b] provide a candidate number, namely 6. As we discussed in Chapter 10, they showed that any given visited page is likely to be one of the last 6 the navigator has seen (although this particular number may be dependent upon the amount of backtracking during the navigation session). Using this number, however, would assume a connection between the likelihood of a page being revisited and its potential contribution to habituation, and the plausibility of this connection is debatable. Note that we might weight the last *n* transitions according to the time spent at the destination, since on an intuitive level more time spent might lead to higher levels of habituation. This approach is practical when dealing with content habituation, but not with navigational habituation, since much of the user's time is likely to be spent scanning content, not on the navigational support. Since click-stream analyses do not provide data on where attention is focused, length of page stay is less informative.

Finally, we might make use of navigation "patches" as the basis for our analysis of habituation in a hyperspace, by simply keeping track of the number of consecutive pages visited within the same patch. This approach could presumably cover both content and navigational volatility, if we assume that the sight architecture is such that, for the most part, related site topics are in fact within the same patches. But again, the user's level of content domain expertise will likely affect her ability to notice these relationships, and therefore her level of habituation. Navigationally, this approach has weight, since for many hierarchically structured sites, patches tend to have similar navigational support (and in fact will only work when pages within a sub-site do in fact have similar navigational support). The trick in using this measurement would be in slicing the site into the appropriate patches of the appropriate sizes, likely to often correspond to sub-sites.

13.4 Volatility predictability

In Chapter 8, we discussed four types of expertise relevant to this thesis: (i) general web navigation expertise, (ii) content domain expertise, (iii) site domain expertise, and (iv) within-site expertise (or expertise regarding a specific Web site due to previous interaction with it). These types of expertise will be important in a discussion of volatility predictability, or the extent to which a navigator is able to anticipate interface changes at the destination page of a transition. Consider the ways in which expertise might help a navigator in making such predictions:

> *General web navigation expertise* allows a user to make predictions based upon knowledge of how Web sites are typically organized and designed. For example, she might predict that by clicking on a hyperlink in a top-level navigational mechanism, the local context hyperlinks along the left

column of the page will change, since many hierarchically organized sites are designed this way.

Content domain expertise allows a user to make predictions based upon knowledge of how information in the domain is structured and interrelated. For example, she might predict that clicking on a “Lost in hyperspace” hyperlink would likely lead to information about disorientation and possibly links to information about navigation design, based on her knowledge of the subject.

Site domain expertise allows a user to make predictions based upon knowledge of how a particular class of sites are typically organized and designed. For example, she might predict that clicking on the name of an author at a bookstore site would lead to a page with a list of books for sale by that author, since many bookstore Web sites are designed this way. (That is, she recognizes an external consistency, as discussed in Chapter 6.)

Within-site expertise allows a user to make predictions based upon previous interaction with a specific Web site. For example, she might predict that within a particular site, clicking on the name of an Italian restaurant will lead to a page with the restaurant’s menu, since other similar hyperlinks within the site have done the same. (That is, she notices both linguistic homogeneity, per Chapter 2, and an internal consistency, per Chapter 6.)

Volatility predictability, in general, relies on a navigator’s ability to recognize hyperlink attributes with predictive power and map them to interface changes from the source page to the destination page. In attempting to measure the concept, we could conceivably make use of any and all such attributes. In hierarchically organized sites, mechanisms tend to have generally homogenous transitional volatility; that is, hyperlinks within the same navigational mechanism tend to result in more or less the same sorts of interface changes, and so one example metric for the predictability level of any given page-to-page transition would be the number of times the user has invoked the same navigational mechanism as the one being used for the transition in question. There are caveats, of course. It is not clear, for example, what the predictability level of a Back button click should be, and in the spirit of DeRose [DeRo89], we might like to treat associative hyperlink lists as “completely unpredictable,” with predictability scores of zero.

13.5 Perceived transitional volatility

Variables based on user exposure to a Web site (volatility habituation and volatility predictability), joined with actual transitional volatility, are proposed to constitute the user’s perceived transitional volatility:

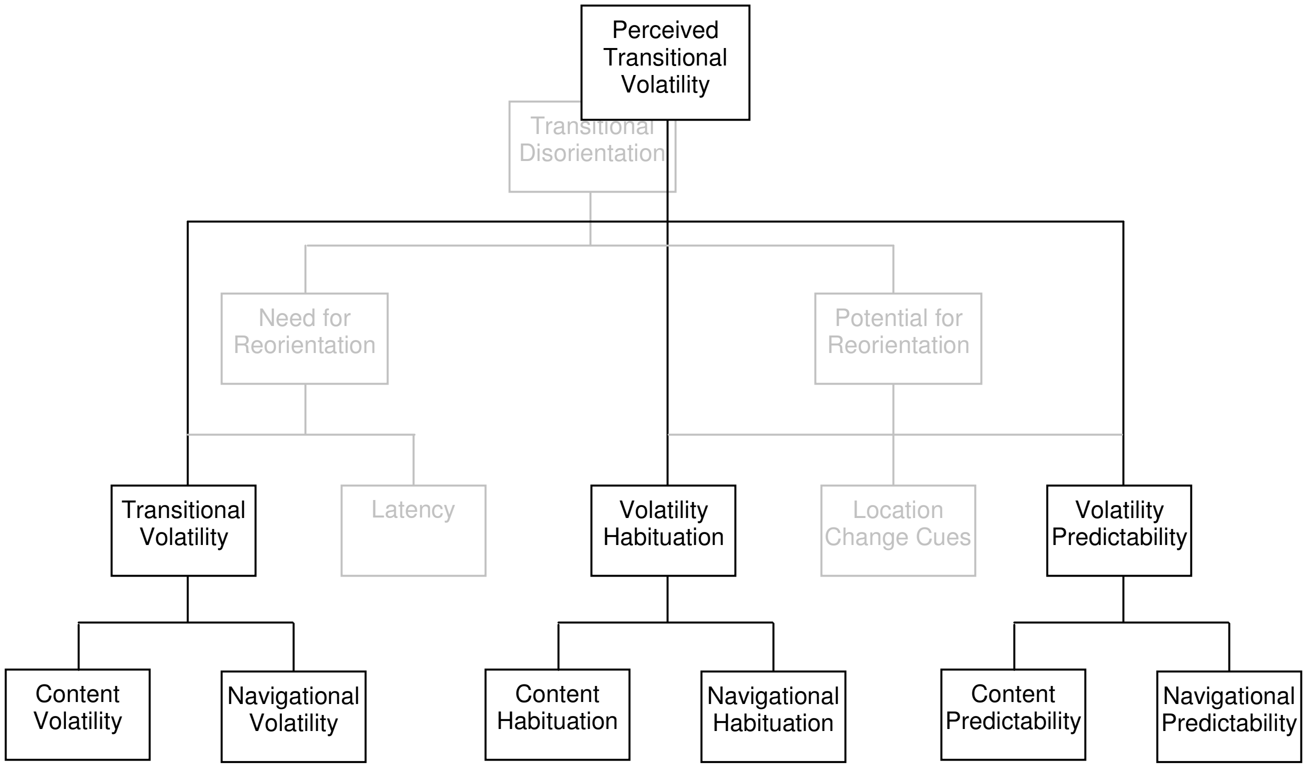

Figure 13.2: Perceived transitional volatility. Actual transitional volatility, volatility habituation, and volatility predictability are proposed as contributing to a user's perception of the extent to which interface changes occur as a result of page-to-page transitions.

More importantly, the user's experience with the Web site should be the cause of any discrepancy between actual and perceived volatility. Habituation should lead to overestimation of actual volatility. Predictability should lead to underestimation of actual volatility.

13.6 Habituate-Predict-Reorient

In the following chapters we will discuss a study motivated by the concepts described in this chapter, and dealing with many of the navigational concepts discussed in the preceding chapters. While these chapters have likely made it abundantly clear that the Web navigation process is extremely complicated, we can sum up with the simple perspective that has been a driving force throughout our discussion:

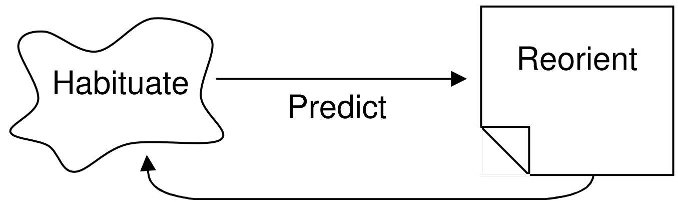

The user becomes habituated within the recent navigation patch. The user predicts content and navigation option changes in page-to-page transitions. The user reorients into the destination page of a transition. The destination page becomes part of the recent navigation patch, continuing the cycle.

PART III: A STUDY OF TRANSITIONAL VOLATILITY

14 Hypotheses

In Parts I and II we examined a space of attributes for the critical players in the Web navigation game, and ways of empirically investigating the interactions between those attributes. The thread throughout our discussion has been a focus on the hyperlink transition as a critical player, and on one of its attributes, namely the interface changes that occur from the source to the destination page, as a potentially fruitful attribute for investigation. In this chapter we will briefly look at four hypotheses that will we be tested by the experiment described in the next chapter.

Recall from Chapter 13 that a volatile navigation session may be assessed in terms of the extent to which new navigation options tend to appear on the destination pages of hyperlink transitions during a navigation session. Our concern here will be with the sorts of effects a more or less volatile navigation session can have.

14.1 H1: Disorientation

H1: A volatile navigation session will cause increased disorientation for a user.

14.2 H2: Site size and complexity

H2: A volatile navigation session will cause a user to perceive a site as being larger and more complex.

14.3 H3: Model of exploration

H3: A volatile navigation session will cause a user to have a less accurate model of her own exploration of a site.

14.4 H4: Perceived global coherence

H4: A volatile navigation session will cause a user to perceive a site as being less globally coherent.

15 Experiment

This chapter describes the experimental design of a study investigating transitional volatility in Web navigation, collecting and analyzing both behavioral and attitudinal data. The study uses the transitional disorientation model discussed in Chapter 13 as the basis for its design and procedure. The model is used and tested as follows:

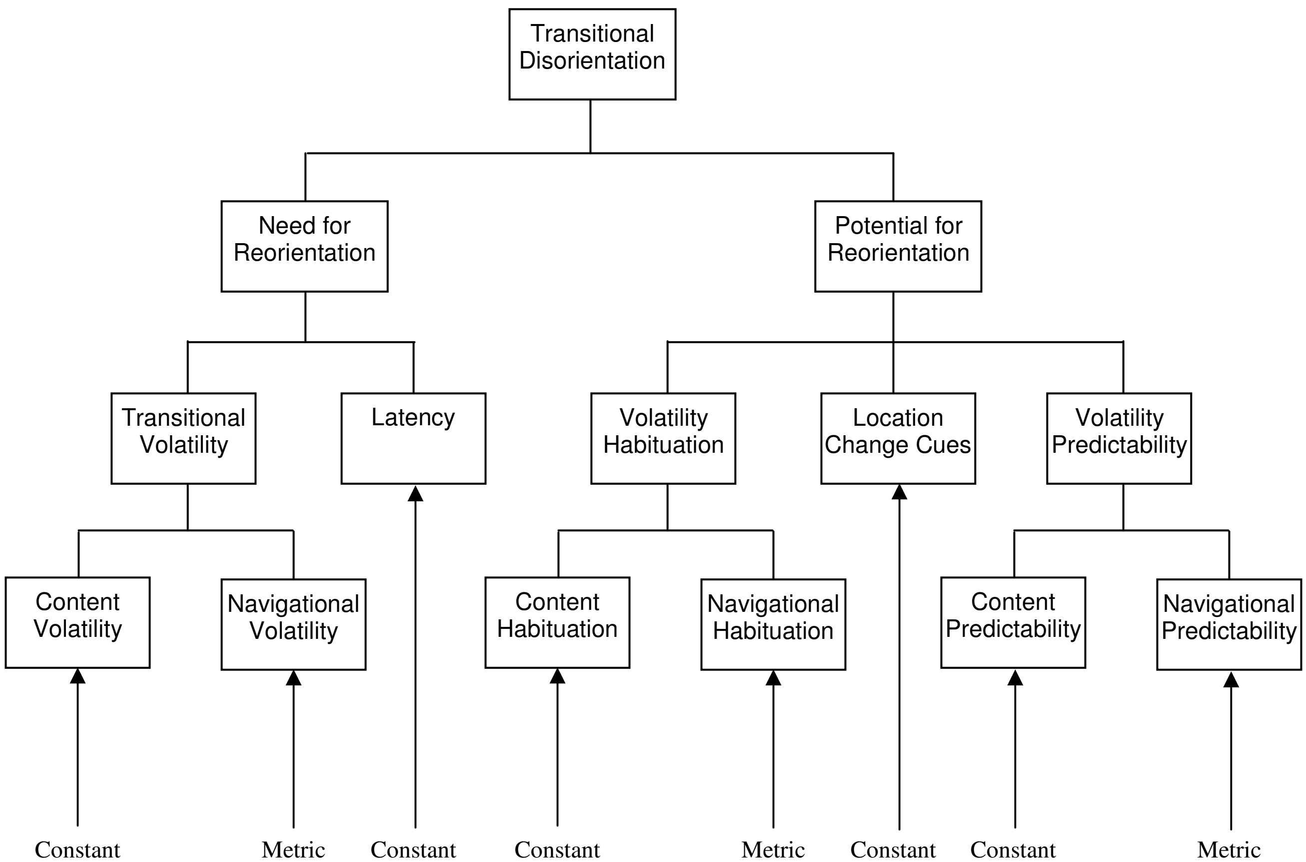

Figure 15.1: Experimental application of transitional disorientation model. Metrics are developed for three factors varying from navigator to navigator: navigational volatility, habituation, and predictability. Other proposed contributing factors are held constant.

Navigational volatility, habituation, and predictability were approximated with metrics, while the other proposed contributing factors were held constant. Content volatility was, more accurately, minimized in the study; there were no font or color changes between pages, and no graphics were used in the study (except on the home page). However, moving from a page with *n* lines of text to a page with *m* lines of text where the change is large enough to present a noticeable difference nonetheless presents some content volatility. Consequently, content habituation and content predictability are more accurately said to be maximized. The design specifications are described in section 15.2. Latency was also, more accurately, minimized, as an effort was made to prevent any page loading time differences between subjects. Stimulus pages were kept on a local machine, and no graphics were used (except on the home page), which is appropriate for low word count pages [Ivor+01].

15.1 Participants

Thirty Stanford University students — half female and half male — participated in the study. The participant age range was 18 to 33 years. Participants were randomly assigned to participate as part of a course requirement. Five females and five males were randomly assigned to one of three experimental conditions, described in the next section. Analyses of age effects were discussed in Chapter 8 and suggest that these participants were likely more efficient navigators than a wider age sample would be.

15.2 Stimuli

The stimulus site used in this study contained text from a government "self-help" legal information site for the California court system. The stimulus site was constructed to control design aspects, as discussed in previous chapters, and contained 100 pages, including a home page. The site was hierarchically structured and organized into five main sections — "Going to Court," "Family Law," "Domestic Violence," "Juvenile Law," and "Small Claims" — each of which contained between 18 and 21 pages and two levels of hierarchical structure. Limiting the maximum hierarchical depth from the home page to three levels may help limit some of the age effects mentioned in section 15.1 [Zaph+00]. The site depth is appropriate given that Web navigators rarely traverse more than two layers (for example, from a top-level page to Level 3 in the hierarchy) before returning to a hub page [Catl+95]. Moreover, the stimulus site follows Chimera and Shneiderman [Chim+94] in that the site structure was "well-formed" down to the third level. Chimera and Shneiderman define this property as being achieved when "each of the two upper levels' items always had subordinates, and no item at the third level had any subordinates." They acknowledge that many real-world overviews are not well-formed, but the use of a well-formed structure may help "account for performance differences attributable to interface design" [Chim+94].

The site contained 19,411 words (including a 46-word site introduction at the home page), with the common progressive increase in amount of content at lower levels of the site hierarchy, as shown in Table 15.1. The word counts place the page lengths in the low to medium word count range in Ivory, Sinha, and Hearst's [Ivor+01] analysis, as discussed in Chapter 5.

Table 15.1 Stimulus site word count analysis
Mean word count for each of the three levels of the stimulus site hierarchy, with progressive word count increase at lower levels.

	Pages	Mean	SD
Level 1 ("Top Level")	5	74.0	51.1
Level 2	19	137.3	107.8
Level 3	75	218.5	116.3

The stimulus site was kept on a local machine to avoid loading time differences, a subject of primary concern in Web navigation [Pik98, Byrn+99b, Niel99]. As mentioned, the constructed site contained no graphics, except on the home page. Diaper and Waeland [Diap+00] reported that graphic content significantly affects information extraction on Web pages by novice users, but not by experienced users. Lack of graphics may therefore limit potential Web expertise effects. Moreover, lack of graphics helped keep loading time and content volatility consistent both within and between subjects.

15.2.1 Conditions

The stimulus site was presented with three different design schemes varying the navigational support provided, each representing a condition:

Full Overview ("FO"): hyperlinks for all pages in the site provided

Partial Overview ("PO"): hyperlinks to the five top-level pages, and to all pages within the sub-site of the current node

Sibling-Child ("SC") (or "Local Context"): hyperlinks to the five top-level pages, and to the siblings and children of the current node

Figure 15.2 shows an example page in each of these conditions:

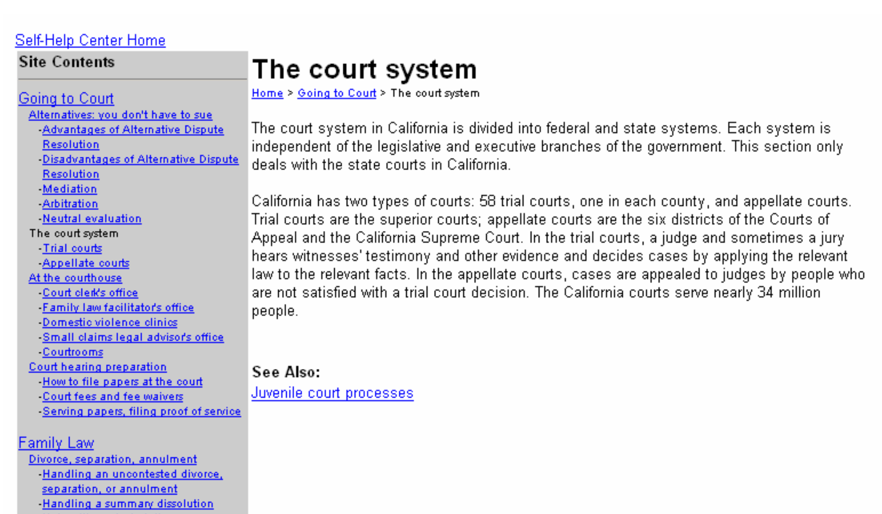

(a)

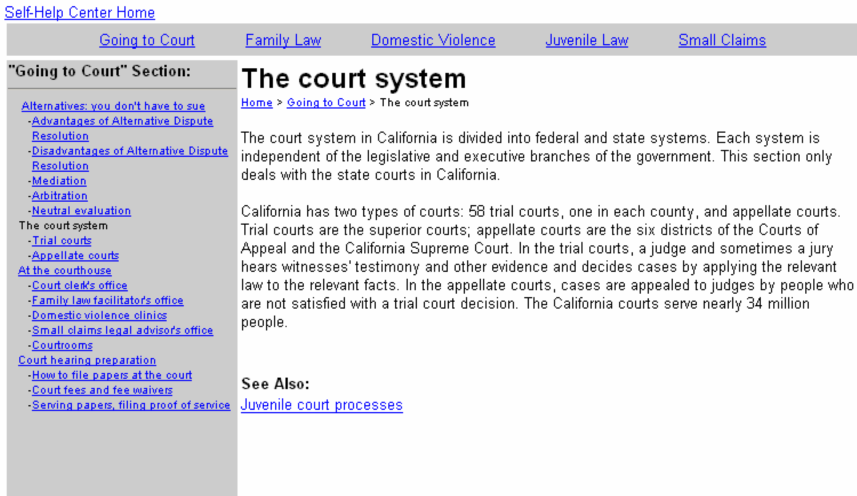

(b)

(c)

Figure 15.2: Experimental condition example. An example page in the (a) Full Overview (FO) condition (only what was above the scroll line for participants is shown), (b) Partial Overview (PO) condition, and (c) Sibling-Child (SC) condition.

The content of the site was exactly the same in all conditions. The content area of the site had the same screen placement for the PO and SC conditions, slightly lower than in the FO condition due to the top-level mechanism across the top of the page. In addition to their unique navigational mechanisms, each condition shared three mechanisms: (i) a "Self-Help Center Home" link located at the top-left corner of the screen, (ii) a breadcrumb list located just below the page title, and (iii) an associative list mechanism located below the body text. The number of associative links varied from page to page, with some pages having none. When there were associative links, they were preceded by "See Also:" in bold text just above them. Level 2 links in these overviews were indented three spaces. Level 3 links were indented three more spaces, and preceded by a hyphen.

The site was presented in 800x600 screen resolution in all conditions. Each condition shared the same font specifications, and an effort was made to apply exactly the same design specifications to all conditions, insofar as the differing design schemes could be treated consistently. All text, both linkable and non-linkable, was presented in Arial/Helvetica font. The page title was in size 18 point font; the body text, associative hyperlinks, and the "Self-Help Center Home" link were in size 10. In the FO condition, top-level links were in size 10. In both the FO and PO conditions, non-top-level links were in size 7.5, as was text in the breadcrumb list mechanism. The full overview, partial overview, sibling list and child list mechanisms were each preceded by a bold title in size 10 and a 1.5 point horizontal line. Each of these navigational mechanisms, in addition to the top-level mechanism in the PO and SC conditions, was presented in boxes with 15% grey shading. Left column mechanisms spanned 200 pixels of screen space, and the content area spanned the remainder of the 600 pixels (with some white space, and room for the browser attributes, including a scroll bar when needed). These specifications led to a line of body text spanning about 75 characters, a reasonable length to support online reading, per Chapter 5.

The appearance of navigational mechanisms as vertical, left-column modules, and their grey, differentiated background color from the page content were appropriate specifications for supporting visual search performance [vanS+01], and follow the common practice of "yellow fever" [Niel00].

The site home page remained the same for each of the three conditions:

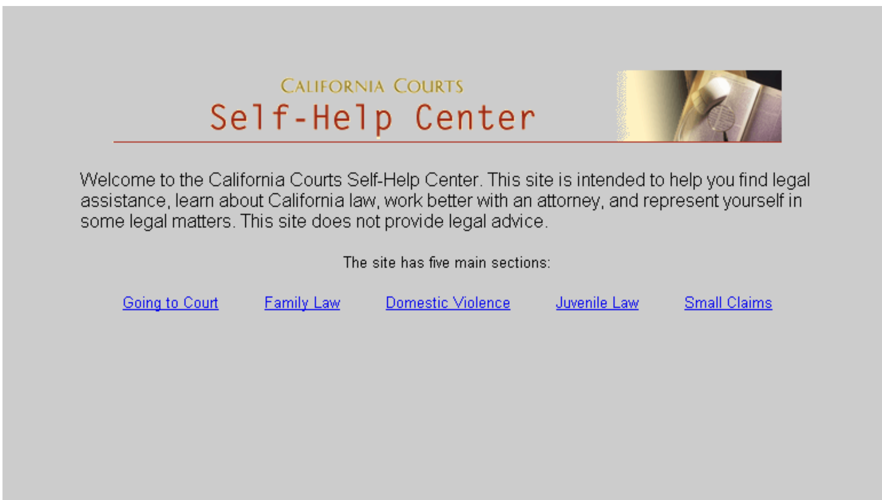

Figure 15.3: Stimulus site home page. The site home page remained the same for each of the three experimental conditions. After completing or abandoning a task, participants reentered the site through this page, and began the navigation session at this page.

Horizontal scrolling was never necessary, given the above specifications. Vertical scrolling was not required for structural navigation support in either the PO or SC conditions, but content did appear below the scroll line for higher word count pages. The first 26 hyperlinks in the full overview mechanism appeared above the scroll line. The full overview, in Chimera and Shneiderman's [Chim+94] classification, per Chapter 3, was in the low-end medium size range. The partial overview was in the small size range, as it fit in one screen display.

15.2.2 Questionnaire

In addition to the stimulus Web site, the study included a navigation questionnaire, composed of five one-page parts. All parts of the navigation questionnaire are available in Appendix C.

Part I asked the participant to provide estimates regarding the size of the site and the extent to which it and each of its main sections had been explored, as well as a few reflective ratings. Part I is available in Appendix C.1.

Part II assessed various usability attributes of the site, using 1-10 agree/disagree Likert scales for statements such as “I generally knew where I was in the Web site.” Numerous rating-focused approaches to assessing aspects of site usability exist [Chin+88, Niel93, Niel94, Smit+97, Frøk+00, Otte+00]. Usability ratings in the research literature are often too general for direct application, but some existing ratings were modified to fit the specific interests of this study. Part II is available in Appendix C.2.

Part III assessed the user’s understanding and model of the condition’s design scheme, and asked the user to provide usability ratings for each of the mechanisms’ and the scheme’s usefulness in finding the desired information during the navigation session. Part III is available in Appendix C.3.

Part IV attempted to assess perceived global coherence of the site by asking the user to rate the relatedness of a sample of 15 page title pairs from the site. Seven of these were pairs for pages with local familial relationships in the site hierarchy: three sibling pairs (one for each hierarchical level), two parent-child pairs (one for L1-L2 and one for L2-L3), and two grandparent-grandchild pairs. Eight were pairs for pages with distal relationships in the hierarchy: four from within the same main section of the site, and four from different main sections. Part IV is available in Appendix C.4.

Part V also attempted to assess perceived global coherence of the site by asking the user to rate the extent to which a sample of nine page titles from the site were central to the site’s main topic, with three pages coming from each of the three hierarchical levels (two pages from four of the five main site sections, and one from the fifth). Part V is available in Appendix C.5.

15.3 Tasks

This study was concerned with directed search, or low complexity fact-finding missions, as discussed in Chapter 7. A set of 25 information-seeking tasks were used in the study, and were randomly ordered. (Participants were not expected to complete all 25 tasks during the navigation session.) Participants received one of two random orderings, five participants in each of the conditions to one ordering, and the remainder to the other. The tasks were typed on 4” by 6” cards, and the task set is available in Appendix D. An example task was “Where can you get information on local department locations and court hours?”

The number of tasks needed to be increased from 15 to 25 due to the potential variability in willingness to abandon a task (an option discussed in the next section); two pilot participants exposed the experimental design problem by completing all 15 tasks in the original set before a 15-minute experimental period had ended (making their data unusable), an issue that did not arise in previous work with similar methodology [Dani01, DaniIP]. In general, the option to abandon a task requires the experimenter to provide enough tasks to ensure that participants will not complete all tasks before the navigation session expires, if data comparison depends upon each doing so.

Because the tasks were presented on physical cards, increasing the number of tasks meant slightly increasing the size of the stack; it is unknown if this had any behavioral or attitudinal effects. In general, the relative tradeoffs between presenting tasks electronically and on physical cards may need further consideration. Presenting the tasks electronically presents further overhead that may be avoided with physical cards. Moreover, physical cards may allow greater flexibility and comfort since the participant may place them anywhere she wishes or hold the card while navigating if comfortable, and need not turn to another screen display to review the current task. (Having the current task on the same computer screen as the stimulus site would mean either using screen real estate, or occluding the task and requiring the participant to shuffle between task and stimulus displays, which, again, presents possibly unnecessary overhead.)

15.4 Procedure

Experimental sessions were separated into three sections. First, participants were given information regarding the general nature of the study, and instructions for participation. The initial background and instructions sheet is available in Appendix A.

Second, participants navigated the stimulus site for 15 minutes, a session length short enough that it was believed to avoid participant fatigue problems [Niel89]. While navigating, screen monitor input was recorded to standard video output for later viewing and coding. This option was preferred, as it was believed to be less invasive than many other recording options. All participants used the same PC, with a wheel scroll mouse, and Internet Explorer was the browser used in this study. The experiment began with the participant clicking on a link to the home page of the stimulus site (and therefore originally entering the site through the front door), and turning over the first 4” by 6” card containing the first task. Participants attempted to complete one information-seeking task at a time, in the provided order, and had the option of abandoning a task if they did not believe they would find the answer. In particular, each participant had been instructed: “Make an effort to answer the questions as best you can, but if you become discouraged and do not believe you will find the answer, you may move on to the next question.” Participants could not later return to previously abandoned tasks. If the participant found the answer to the information-seeking task, she needed only to highlight the information, as though to copy and paste it to consult the information later. (Since the screen output was recorded, it was not necessary for the participant to provide a physical record of completed tasks.) This experimental design decision was meant to minimize the amount of time spent recording the extracted answer, as the concern of this study was on the navigation process. After participants successfully completed tasks, they clicked on an "ANSWER QUESTION" link in the Favorites bar of the browser to stamp the completion, and returned to site via the home page to begin the next question; similarly, when they abandoned a task, they clicked on a "SKIP QUESTION" link, also in the Favorites bar of the browser, and returned to the site via the home page. The procedure of returning to the home page after each task is methodologically similar to Otter and Johnson [Otte+00], who point out that such practice ensures that all participants start tasks from the same point in the site (although, of course, one cannot control the navigation histories of the different participants at the start of each task).

Thirdly, participants completed the post-navigation questionnaire materials. This section lasted about 20 minutes. Although there was no time pressure, participants did not vary widely in the amount of time spent on the questionnaire, and are not believed to have varied widely in the amount of thinking devoted to the post-navigation questions and ratings.

Session and post-navigation data together assessed efficiency, effectiveness, and subjective satisfaction, as one cannot assume these to be correlated [Niel+94, Frøk+00].

15.5 Behavioral data collection

Navigational click-stream data was collected for analysis from the recorded screen output. Page visits and time of arrival (in seconds, starting from zero at the beginning of the navigation session) were recorded. For each navigational act, the following attributes of the page visit were recorded:

(i) Node code, which uniquely identified the page visited in the site hierarchy. Each node code indicated the sibling order of the top-level category in which the page appeared, and, if applicable, the sibling order of the second and/or third level categories.

(ii) The navigational mechanism used to arrive at the destination page (i.e. Full Overview, Partial Overview, Top-Level Links, Sibling List, Child List, Breadcrumb List, Associative List, Back button, etc.).

(iii) Time of arrival.

Task completion and abandonment and site reentries via the home page were also recorded, but were treated as separate from other navigational acts. The data collection allowed for frequencies and proportions of the following to be computed for each individual participant:

Completed tasks
Abandoned tasks

Navigational actions
Mean time per page
Unique page visits
Page revisits
Home page visits

Top-level page visits
Intermediate-level page visits
Bottom-level page visits
Top-level switches

Top-level time spent
Intermediate-level time spent
Bottom-level time spent

Hyperlink clicks
Back button, History List
Forward button clicks

Sibling transitions
Child transitions
Parent transitions
Grandchild transitions
Grandparent transitions

Distal transitions

Upward movements
Downward movements
Lateral movements

Top-Level mechanism clicks
Full overview clicks
Partial overview clicks
Sibling list clicks
Child list clicks
Breadcrumb list clicks
Associative list clicks

A "top-level switch" was defined similarly to previous work [DaniIP], namely navigational movements such that the source and destination pages were in different main sub-sites. The methodology also allowed for the application of a number of click-stream metrics, analyzed in the next chapter.

16 Results

This chapter reports the complete results of the study described in the previous chapter, with analyses of the behavioral and attitudinal data collected. A total of 1730 navigational acts were recorded. The 30 participants visited a total of 2400 pages, including task completions and site reentries via the home page, which were treated separately from other intra-site navigational transitions. Three ratings (for individual participants and individual questions) in the post-navigation questionnaire were left blank; in each case the mean rating for the remaining nine participants in that participant's condition was used.

The chapter will begin with an overview of the navigational volatility, habituation, and predictability metrics, followed by between-condition analyses. Finally, effects of the metrics used will be reported.

16.1 Metrics

For the Partial Overview and Sibling-Child (Local Context) conditions, each participant's click-stream data resulted in indexes for (i) navigational volatility, (ii) navigational habituation, and (iii) navigational predictability. The hyperlinks stayed the same for all navigational acts in the Full Overview condition, and so the same analyses were not appropriate. The Full Overview condition served in some cases as a control, used in interpreting results, as we will discuss in the next chapter.

For each transition, navigational volatility was defined as the number of hyperlinks appearing on the destination page that did not appear in the same screen location as on the source page. The participant's volatility score was the mean navigational volatility across all navigational acts during the experimental session.

For each transition, navigational habituation was defined as the number of previous consecutive pages the participant had visited within the same main sub-site; that is, habituation was based upon how long (in terms of page visits) the navigator had remained within the same information-seeking patch (Recall that this metric works well when the stimulus site is such that pages within the same sub-sites have similar navigational support, which is common in hierarchically organized sites, and was the case in this study.) The participant's habituation score was the mean navigational habituation across all navigational acts during the experimental session.

For each transition, navigational predictability was defined as the number of times the user had previously used the same navigational mechanism, during the navigation session, as the mechanism being used for the transition in question; that is, predictability was based upon the participants' previous usage of the current mechanism being used for a hyperlink transition. The participant's predictability score was the mean navigational predictability across all navigational acts during the experimental session.

Note a few details. When participants completed a question, they reentered the site via the home page, and in these cases it was sometimes the case that a user would reenter into the same main sub-site as before the task completion. However, each time a task was completed, the navigational habituation of the next transition was still measured at zero (the start of a new habituation patch). Note that regardless of the number of previous times a user invoked an associative hyperlink, the predictability of a such a transition was always given a zero value, in the spirit of the commonly agreed upon unpredictable nature of associative hyperlinks. Back button clicks in this study were not assigned a predictability score, as any such score was difficult to justify. Future similar work might attempt to develop and empirically validate a predictability score for Back button clicks that accounts for their likely somewhat special nature.

Across the Partial Overview and Local Context conditions, the mean navigational volatility score was 4.06 (SD = 1.11, Median = 4.27, Min = 1.51, Max = 5.72). The mean navigational habituation score was 1.43 (SD = 0.86, Median = 1.16, Min = 0.76, Max = 4.11). The mean navigational predictability score was 12.30 (SD = 5.82, Median = 12.67, Min = 3.25, Max = 22.90). The two conditions were not significantly different in their navigational volatility scores, but were for both habituation and predictability. Partial Overview participants had a mean habituation score of 1.00 (SD = 0.32), while Local Context participants had a mean score of 1.85 (SD = 1.03), significantly higher (t-test, p = 0.031). Thus, Local Context participants tended to stay within navigational patches more than Partial Overview participants.

Local Context participants had a mean predictability score of 7.89 (SD = 3.34), while Partial Overview participants had a mean score of 16.71 (SD = 4.13), significantly higher (t-test, p = 0.000). Thus, Partial Overview participants tended to use navigational mechanisms repeatedly more than Local Context participants, which is not surprising since the navigational scheme in that condition was composed of one less mechanism than that of the Local Context condition.

16.2 Click-stream analysis

Click-stream data in this study included the destination page with time stamp, as is standard in such analyses, as well as the specific mechanism used, the hierarchical source-destination relationship, and the direction of movement for each hyperlink transition.

16.2.1 Page visits

Each Web page in the stimulus site had a unique code, which included its hierarchical level. Attention was paid in data analysis to whether the page was being visited for the first time or revisited and the hierarchical level of the page visit. The amount of time spent at the home page and each of the hierarchical levels was also analyzed (with the full navigation session lasting 900 seconds). Table 16.1 summarizes the page visitation data for each of the three conditions:

Table 16.1 Page Visit Data
Frequencies and proportions of unique page visits and revisits, and hierarchical level of page visits, with 1-way ANOVA p-values. Time measurements are in seconds.

	FO		*PO*		*SC*		*1-way ANOVA*
	Mean	SD	Mean	SD	Mean	SD	p
Page Visits	63.2	12.9	91.8	27.0	85.0	29.8	0.036 *
Unique	31.3	8.0	44.6	9.3	33.0	9.2	0.004 **
Revisit	31.9	7.6	47.2	19.1	52.0	22.8	0.044 *
Unique %	49.46%	7.65%	50.31%	8.52%	40.38%	9.05%	0.024 *
Revisit %	50.54%	7.65%	49.69%	8.52%	59.62%	9.05%	0.024 *
Home	12.1	2.0	11.5	3.7	13.0	3.4	0.567
Level 1	13.9	2.8	23.3	9.1	23.4	7.5	0.007 **
Level 2	13.9	4.7	18.6	5.2	29.2	14.2	0.003 **
Level 3	23.3	6.4	38.4	12.2	19.4	9.2	0.000 **
Home %	19.52%	3.05%	12.59%	2.05%	16.51%	5.92%	0.003 **
Level 1 %	22.31%	3.57%	24.73%	5.16%	27.90%	4.53%	0.032 *
Level 2 %	21.65%	4.31%	20.88%	4.83%	33.02%	6.59%	0.000 **
Level 3 %	36.53%	6.12%	41.80%	3.65%	22.57%	5.30%	0.000 **
Home Time	117.3	22.6	97.7	24.5	142.5	25.5	0.001 **
Level 1 Time	233.0	55.9	197.3	48.5	181.2	36.0	0.061
Level 2 Time	206.5	43.8	211.7	51.2	324.6	66.2	0.000 **
Level 3 Time	343.2	86.9	393.3	31.2	251.7	49.1	0.000 **
Home Time %	13.04%	2.51%	10.86%	2.72%	15.83%	2.83%	0.001 **
Level 1 Time %	25.89%	6.21%	21.92%	5.40%	20.13%	4.00%	0.061
Level 2 Time %	22.95%	4.87%	23.52%	5.69%	36.07%	7.36%	0.000 **
Level 3 Time %	38.13%	9.65%	43.70%	3.47%	27.97%	5.46%	0.000 **

*Significant at p < 0.05 **Significant at p < 0.01

The mean revisitation rate for the 30 participants across the three conditions was 53.28% (SD = 9.33%).

16.2.2 Method of access

A total of 407 navigational acts in the Full Overview condition, 709 in the Partial Overview condition, and 614 in the Sibling-Child (Local Context) condition were recorded. Table 16.2 summarizes the method of access data for these navigational acts, which includes the navigational mechanisms used in hyperlink transitions:

Table 16.2 Method of Access Data
Frequencies and proportions of the navigational mechanisms used to access pages, with 1-way ANOVA p-values.

	FO		*PO*		*SC*		*1-way ANOVA*
	Mean	SD	Mean	SD	Mean	SD	p
Answer/Skip Question	11.1	1.5	9.7	3.3	11.2	3.2	0.424
Site Reentry	11.4	2.1	11.2	3.6	12.4	3.1	0.638
Navigational Acts	40.7	11.6	70.9	23.3	61.4	26.5	0.012 *
Home Page Link	0.2	0.4	0.6	0.7	1.0	1.5	0.209
Breadcrumb	0.3	0.7	0.7	1.5	1.0	1.4	0.464
Associative List	4.9	4.4	9.0	5.1	7.3	5.0	0.183
Top-Level	-	-	11.0	6.5	5.6	4.7	0.047 *
Full Overview	33.0	10.0	-	-	-	-	-
Partial Overview	-	-	46.3	12.4	-	-	-
Sibling List	-	-	-	-	11.4	7.1	-
Child List	-	-	-	-	25.5	11.0	-
Back	2.1	2.7	3.3	3.6	9.6	6.3	0.002 **
Forward	0.1	0.3	0.0	0.0	0.0	0.0	0.381
History List	0.1	0.3	0.0	0.0	0.0	0.0	0.381
Home Page Link %	0.46%	1.00%	0.80%	1.01%	1.63%	2.40%	0.269
Breadcrumb %	0.79%	1.94%	0.80%	1.71%	1.44%	1.58%	0.643
Associative List %	11.38%	9.52%	12.18%	5.19%	11.77%	5.32%	0.967
Top-Level %	-	-	14.37%	6.92%	8.03%	6.08%	0.043 *
Full Overview %	82.18%	12.29%	-	-	-	-	-
Partial Overview %	-	-	67.27%	11.80%	-	-	-
Sibling List %	-	-	-	-	18.41%	7.89%	-
Child List %	-	-	-	-	42.68%	9.26%	-
Back %	4.82%	5.96%	4.57%	4.79%	16.06%	8.38%	0.001 **
Forward %	0.19%	0.60%	0.00%	0.00%	0.00%	0.00%	0.381
History List %	0.19%	0.60%	0.00%	0.00%	0.00%	0.00%	0.381

*Significant at $p < 0.05$ **Significant at $p < 0.01$

16.2.3 Relationship and direction

Because the stimulus site was hierarchically organized, the hierarchical relationship of a source-destination transition (for example, to a parent destination, or to a sibling destination) could be recorded. "Distal" movements were to destination pages outside the immediate family of the source page. Top-level switches were defined as movements between main sub-sites — that is, inter-sub-site transitions. Additionally, the direction of movement could be recorded; moving from a page at the third hierarchical level of the site to a top-level page, for example, would be an upward movement. Table 16.3 summarizes the relationship and direction of movement data:

Table 16.3 Relationship and Direction of Movement
Frequencies and proportions of hierarchical relationship and direction of page-to-page transitions, with 1-way ANOVA p-values.

	FO		*PO*		*SC*		*1-way ANOVA*
	Mean	SD	Mean	SD	Mean	SD	p
Top-Level Switches	10.1	3.8	17.9	10.4	14.1	7.2	0.093
Navigational Acts	40.7	11.6	70.9	23.3	61.4	26.5	0.012 *
Sibling	6.5	3.4	21.4	10.1	15.7	9.8	0.002 **
Child	8.5	3.9	13.0	2.4	25.6	10.9	0.000 **
Parent	3.7	2.2	4.4	4.0	7.1	4.5	0.115
Grandchild	4.3	1.6	6.6	3.0	0.1	0.3	0.000 **
Grandparent	0.3	0.5	0.5	0.7	1.1	1.1	0.086
Distal	17.4	7.2	25.0	10.1	11.7	7.2	0.005 **
Sibling %	15.44%	5.14%	29.06%	7.57%	24.54%	9.83%	0.002 **
Child %	21.27%	7.36%	19.91%	6.62%	42.84%	9.20%	0.000 **
Parent %	9.20%	4.74%	6.03%	3.76%	12.32%	7.36%	0.054
Grandchild %	11.15%	5.13%	9.42%	3.23%	0.09%	0.30%	0.000 **
Grandparent %	0.62%	1.01%	0.72%	1.05%	1.55%	1.43%	0.172
Distal %	42.32%	10.69%	34.87%	5.66%	18.28%	8.48%	0.000 **
Up	8.9	4.0	16.0	7.3	13.5	6.2	0.040 *
Down	20.4	5.1	25.6	6.2	29.9	13.4	0.079
Lateral	11.4	5.4	29.3	12.5	17.9	10.1	0.001 **
Up %	21.42%	5.01%	22.11%	4.37%	22.05%	5.46%	0.943
Down %	51.60%	9.33%	37.46%	6.83%	49.57%	8.86%	0.002 **
Lateral %	26.98%	9.15%	40.44%	8.61%	28.00%	10.14%	0.005 **

*Significant at $p < 0.05$ **Significant at $p < 0.01$

16.2.4 Duration of page visit

As shown in Table 16.1, participants with a Full Overview tended to visit fewer pages during the 15-minute navigation session. So, they tended to spend more time on any given Web page. Full Overview participants averaged 14.86 seconds per page (SD = 3.45), Partial Overview participants averaged 10.68 seconds (SD = 3.43), and Local Context participants averaged 11.76 seconds (SD = 3.94). A 1-way ANOVA showed a significant effect of condition on mean time per page (p = 0.041).

Transition attributes were also looked to for data on time spent on a page. Across conditions, there was a significant effect of hierarchical direction of movement on duration of page visit, with downward movements resulting in the most time spent at the destination page (1-way ANOVA, p = 0.000). Similarly, navigators spent more time at pages deeper in the site hierarchy (1-way ANOVA, p = 0.000). Both of these findings are linked to the effect of page word count on duration of visit, as summarized in Table 16.4. Positive correlations between a given page's word count and mean time spent on at that page were found regardless of the hierarchical level of the page and regardless of the experimental condition (99 of the 100 pages in the stimulus site were included in the analyses, excluding the home page since participants returned to the home page to prepare for the next task by reading the question card, affecting the amount of time spent there. Note that in some cases, a Web page was never visited by any of the 10 participants within a condition.

Table 16.4 Word count and duration of visit
Correlations between page word count and mean time spent, across all page visits and within hierarchical level and condition.

	Pages	*r*
Across Level/Condition	95	0.465
Level 1	5	0.641
Level 2	19	0.746
Level 3	70	0.403
FO	75	0.362
PO	89	0.429
SC	81	0.434

Across conditions, non-Back button transitions resulted in a mean time spent at the destination page of 11.70 seconds (SD = 11.30), and Back button clicks resulted in a mean of 5.67 seconds (SD = 6.99), significantly lower (t-test, p = 0.000). Distal and non-distal movements in the site hierarchy did not result in significant differences in time spent at the destination. Top-level switches were also not predictive of duration of page stay. Moreover, across the Partial Overview and Local Context conditions (which had volatility and predictability scores), neither navigational volatility nor navigational predictability of a transition correlated with the amount of time spent at the destination page.

16.2.5 List position bias

Each page visited within the site had a list position in the full overview mechanism. As mentioned in the previous chapter, the first 26 of these hyperlinks were above the scroll line, with the remainder requiring vertical scrolling to be viewed in the Full Overview condition. The list position of a hyperlink in this overview mechanism was defined by its height in the list; the top-most hyperlink had a list position of 1, the link below it a list position of 2, and so on. A 1-way ANOVA showed a significant effect of condition on the mean overview position of visited pages, with Full Overview condition participants invoking links higher up in the overview (p = 0.044). Full Overview participants had a mean list position index for their visited pages of 30.32 (SD = 4.55), Partial Overview participants had a mean list position index of 35.37 (SD = 4.83), and Local Context participants had a mean list position index of 36.36 (SD = 6.75).

16.3 Usability and Performance

Performance was not a primary concern of this study, as it has been looked at extensively in past work, but some data is worth noting. In general, the navigational schemes subjectively showed equal levels of usefulness. Table 16.5 summarizes the usefulness ratings of the navigational scheme and individual navigational mechanisms available to participants in each of the three conditions:

Table 16.5 Navigational Scheme and Mechanism Usefulness Ratings
Participants rated the usefulness of each available navigational mechanism in information-seeking during the experimental navigation session.

	FO		PO		SC		1-way ANOVA
	Mean	SD	Mean	SD	Mean	SD	p
All Navigation Options	5.7	1.3	5.4	2.5	5.7	2.6	0.940
Home Page Link	2.1	1.4	4.0	3.7	4.8	3.2	0.120
Breadcrumb List	3.8	2.9	4.0	1.8	4.8	3.0	0.675
Associative List	4.2	2.0	3.8	3.2	3.9	2.6	0.940
Top-Level Links	-	-	4.8	2.9	3.8	2.3	0.391
Full Overview	8.3	0.9	-	-	-	-	-
Partial Overview	-	-	7.0	3.0	-	-	-
Sibling List	-	-	-	-	7.1	1.4	-
Child List	-	-	-	-	7.4	2.6	-

Mechanisms presented in the left column of the screen (Full Overview, Partial Overview, and Sibling and Child Lists for the three respective conditions), were consistently rated more useful than other available mechanisms, including the top-level link mechanisms in the Partial Overview and Local Context conditions.

Partial Overview condition participants abandoned fewer information-seeking tasks than those with local context support only. Local Context participants abandoned 3.1 tasks on average (SD = 2.1), while Partial Overview participants averaged 1.4 (SD = 1.2), significantly lower (t-test, p = 0.041). However, Full Overview participants, with a mean of 1.6 abandoned tasks (SD = 2.0) did not show a significant difference (t-test, p = 0.115).

Note that, across the Partial Overview and Local Context conditions, the revisitation rate of a user appeared to be a bad sign all around, correlating with a tendency to abandon information-seeking tasks (r = 0.505, p = 0.023). Revisitation rate similarly correlated negatively with "The site was easy to navigate" (r = -0.537, p = 0.015) and "It was easy to find the information I needed in the site" (r = -0.498, p = 0.025). (Although proportion of Back button clicks and revisitation rate are intimately related, Back button use did not provide the same predictive power.)

16.4 Mental Models

Participants estimated the site's size and their own exploration of both the site and of its main sub-sites, additionally giving subjective size ratings. Table 16.6 summarizes these perceptions:

Table 16.6 Site Perception and Exploration Models
Estimate (with accuracy) of site size, exploration of site, exploration of its main sub-sites, and site perception ratings, with 1-way ANOVA p-values.

	FO		*PO*		*SC*		*1-way ANOVA*
	Mean	SD	Mean	SD	Mean	SD	p
Site Size Estimate	80.0	23.0	109.5	62.4	57.0	33.6	0.037 *
Estimate Error	20.0	23.0	45.5	41.2	48.0	25.0	0.098
"The site seemed large" Rating	7.6	2.2	8.3	1.6	5.7	2.7	0.035 *
"The number of navigation options (links) was overwhelming" Rating	6.3	2.1	7.5	1.7	4.3	2.4	0.007 **
Unique Page Visit Estimate	26.4	12.7	36.5	12.7	27.7	10.6	0.145
Actual Unique Page Visits	31.3	8.0	44.6	9.3	33.0	9.2	0.004 **
Estimate Error	7.5	4.4	9.7	7.0	10.5	7.2	0.557
Estimate Error (percentage)	40.47%	42.71%	30.79%	24.01%	45.61%	40.11%	0.659
Section Exploration Error	13.9	4.9	19.7	8.7	17.1	9.6	0.283

*Significant at $p < 0.05$ **Significant at $p < 0.01$

Presented with local context navigation only, participants viewed the site as significantly smaller. Participants in the three conditions did not differ in their ability to accurately estimate the extent to which they had explored the site or its main sub-sites.

There was a significant effect of condition on perceived global coherence of the stimulus site, with Local Context participants viewing the site as more tightly related (1-way ANOVA, p = 0.027). Local Context participants had a mean perceived global coherence index of 7.35 (SD = 0.93), while Partial Overview participants had a mean index of 6.54 (SD = 0.98), and Full Overview participants had a mean index of 6.19 (SD = 0.87).

Finally, across the Partial Overview and Local Context conditions, the section exploration error (the extent to which a user understood her own exploration of the main sub-sites) correlated positively with the proportion of lateral movements in the space (r = 0.445, p = 0.049), as well as the total number of unique page visits (r = 0.450, p = 0.046).

16.5 Transitional effects

This section summarizes the relationships between navigational volatility, habituation, and predictability found in this study. Analyses were done both within and between conditions, to account for potential effects of the navigational schemes of the differing conditions.

16.5.1 Navigational volatility

Navigational volatility correlates in the Partial Overview condition are shown in Table 16.7:

Table 16.7 Navigational Volatility Correlates, Partial Overview Condition
Transitional volatility correlations with disorientation, site size perception, perceived global coherence, the site exploration model, and perceived ease of navigation and site organization for Partial Overview (PO) condition participants. The statements in quotes were rated on a 1-10 agree/disagree scale.

Disorientation	*r*
"When I felt lost, it was easy to reorient myself in the site"	-0.910
"I generally knew where I was in the Web site"	-0.821
"I felt lost and needed to reorient myself"	0.599
Site Size Perception	*r*
"The site seemed large"	0.515
"The number of navigation options (links) was overwhelming"	0.333
Perceived Global Coherence	*r*
"The information in the site seemed to be tied together and well connected"	-0.843
Relatedness ratings for "distal" site topics	-0.647
Perceived global coherence metric (mean topic relatedness rating)	-0.529
"The topics in the site were related to one another and pertained to one coherent topic"	-0.398
Exploration Model	*r*
Section exploration error	0.452
Ease of Navigation and Site Organization	*r*
"The site was easy to navigate"	-0.833
"The information in the site was well organized"	-0.562
"The navigation options (links) were well organized on the page"	-0.493
"It was easy to return to pages I had previously visited"	-0.355
Rated usefulness of navigation options	-0.313

Within the condition, navigational volatility had significant effects on disorientation levels ($p = 0.000$), perceived global coherence ($p = 0.002$), and ease of navigation ($p = 0.003$). A highly volatile navigation session predicted increased disorientation, decreased perceived global coherence, and decreased ease of navigation. Its effect on site size perception and accuracy of the user's exploration model within the Partial Overview condition is not conclusive; more measures will be needed in future work in order to determine whether there is mounting evidence of an effect. Although predictive of subjective site size reports, navigational volatility was not predictive of the actual site size estimates of Partial Overview condition participants. Somewhat surprisingly, navigational volatility also negatively correlated ($r = -0.806$) with user ratings for, "It was easy to find the information I needed in the site," ($p = 0.005$), presumably a measure of perceived information-seeking success.

Navigational volatility correlates in the Local Context condition are shown in Table 16.8:

Table 16.8 Navigational Volatility Correlates, Sibling-Child (Local Context) Condition
Transitional volatility correlations with site size perception, perceived global coherence, and the site exploration model of Sibling-Child (SC) condition participants. The statements in quotes were rated on a 1-10 agree/disagree scale.

Site Size Perception	*r*
"The site seemed large"	0.655
"The number of navigation options (links) was overwhelming"	0.351
Site size estimate (in number of pages)	0.331
Perceived Global Coherence	*r*
Centrality of top-level site topic ratings	0.799
Centrality of site topic ratings (mean for sample)	0.475
"The topics in the site were related to one another and pertained to one coherent topic"	0.471
Perceived global coherence metric (mean topic relatedness rating)	0.449
"The information in the site seemed to be tied together and well connected"	0.378
Exploration Model	*r*
Section exploration error	0.370

Within the condition, navigational volatility had significant effects on site size perception (p = 0.040) and perceived global coherence (p = 0.006). Note, however, that the extent to which centrality ratings measure perceived global coherence is unclear (discussed in the next chapter). If we choose not to rely heavily or solely on such measures (as might be reasonable), a sign test for the six measures (including relatedness ratings for "distal" site topics) shows a significant positive relationship between navigational volatility and perceived global coherence (p = 0.016). A highly volatile navigation session predicted increased site size perception and increased perceived global coherence of the site. Its effect on the accuracy of the user's exploration model within the Local Context condition, as with the Partial Overview condition, is not conclusive.

Across both the Partial Overview and Local Context conditions, increased navigational volatility predicted decreased reported ease of navigation, correlating negatively with "The site was easy to navigate" (r = -0.574, p = 0.008), and increased disorientation, correlating negatively with "I generally knew where I was in the Web site" (r = -0.533, p = 0.016), negatively and approaching significance with "When I felt lost, it was easy to reorient myself" (r = -0.411, p = 0.072), and positively, but not approaching significance, with "I felt lost and needed to reorient myself" (r = 0.318).

16.5.2 Habituation and predictability

Users provided an agree/disagree rating for "The navigation options (links) seemed to stay the same as I navigated," assessing perceived navigational volatility. In the Local Context condition, the regression equation with predictors being actual navigational volatility, navigational habituation, and navigational predictability, was: Lack of Volatility(Perceived) = 8.90 – 1.58*Volatility(Actual) – 1.06*Habituation + 0.645*Predictability (R-sq = 35.8%). In the Partial Overview condition, the regression equation, with the same predictors, was: Lack of Volatility(Perceived) = 14.30 – 1.33*Volatility(Actual) – 2.69*Habituation – 0.068*Predictability (R-sq = 53.3%).

The extent to which a navigator tended to have remained within a navigational patch at the time of a page-to-page transition was generally not predictive of attitudinal data or site perceptions. Across the Partial Overview and Local Context conditions, the navigational habituation metric correlated positively with "It was easy to understand where in the site clicking on a link would take me" (r = 0.452, p = 0.046).

A high level of predictability during the navigation session significantly predicted a decreased proportion of Back button clicks (r = -0.678, p = 0.001). Note also that the user's reported perception of navigational predictability, measured by "It was easy to understand how the page would change when I clicked on a link," approached significance in negatively correlating with the tendency to abandon a task (r = -0.423, p = 0.063).

Finally, note that perceived (lack of) navigational volatility appears at least promising in capturing the notion of predictability during the navigation session. The user rating for perceived lack of navigational volatility, "The navigation options (links) seemed to stay the same as I navigated," correlated positively with each of the three predictability reports by participants. The relationship was significant with "It was easy to predict where in the site clicking on a link would take me" (destination prediction) (r = 0.559, p = 0.010), approached significance with "The (underlined) link titles well summarized the information they led to" (information prediction) (r = 0.407, p = 0.075), and was positive (but not significant) with "It was easy to understand how the page would change when I clicked on a link" (transitional volatility prediction) (r = 0.328). Perceived lack of navigational volatility also correlated negatively, but not significantly, with a presumed (and validated) sign of low predictability, Back button usage (r = -0.347).

17 Discussion

We will now discuss the results reported in the previous chapter, beginning with how they inform the hypotheses proposed in Chapter 14. Although the other specific click-stream data was not the focus of this study, it will be useful to briefly discuss them in relation to previous work in Web navigation behavior.

17.1 Hypotheses

The four hypotheses from Chapter 14 dealt with the potential relationship between a transitionally volatile navigation session and (i) disorientation, (ii) perceived site size and complexity, (iii) a user's model of site exploration, and (iv) perceived global coherence.

17.1.1 Disorientation

H1: A volatile navigation session will cause increased disorientation for a user.

Across the Partial Overview and Local Context conditions, a volatile navigation session and reported disorientation correlated negatively. Within the Partial Overview condition, but not within the Local Context condition, the relationship was especially strong. The evidence appears to be supportive of H1, but a more detailed discussion is needed.

First and foremost, recognize that in the Partial Overview condition, the extent to which the user is confronted with new navigation options as she moves from one page to another is heavily influenced by how frequently she makes "top-level switches," or movements between pages within different main sub-sites. In order to understand the probable causal relationships, we must consider these movements. H1 suggests that a volatile navigation session is a causal influence on disorientation. But we must consider another story:

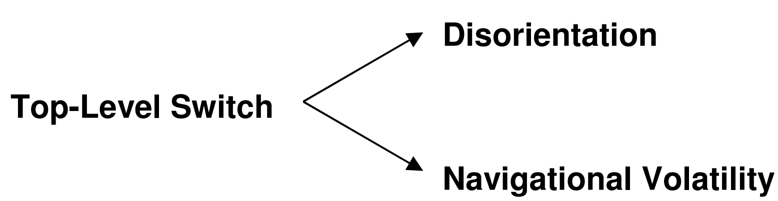

We know that "Top-Level Switch → Navigational Volatility" is true just by the definition of the two terms (and consequence of the design schemes). In the above model, top-level switches cause both disorientation and, as a by-product (and consequence of the design scheme), navigational changes in the Web interface. We might extend this model to include a causal influence of disorientation on top-level switches, creating a potentially ugly cycle:

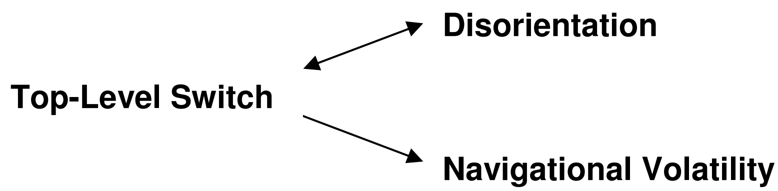

This additional causal relationship is at least plausible; top-level nodes are reorienting nodes, and may be viewed that way by navigators. Moreover, recall that when scent within a local information patch appears low, navigators are likely to move to another patch (which may involve making a top-level switch). A cycle like this might explain why the correlations between disorientation and navigational volatility were so high; the disoriented user seeks reorientation in top-level switches, which cause more disorientation, leading to more top-level switches, and all the while the cycle includes increased navigational volatility.

In saying that top-level switches cause disorientation, we ought to be able to say what it is *about* a top-level switch, as opposed to transitions within the same sub-site, that might give such transitions their causal power. The experimental design, described in Chapter 15, presented the attributes of the stimulus site. The site was stripped down in such a way as to eliminate or minimize any differences between inter- and intra-sub-site transitions, other than their differences in navigational volatility, as determined by the navigational scheme of the condition. They differ in one other way: the extent to which the specific information in the source and destination pages of the transition is likely to be related. That is, a non-visual change occurs: the user moves from one kind of information to another which is relatively unrelated. Although lostness appears to occur independently of a navigator's content domain expertise [Elm+85], disorientation nonetheless could be an effect of moving from one kind of information (assuming the user has some grasp of the sort of information space she is in) to another (assuming she notices on some level, by the specific content, that she is not in Kansas anymore). That is a story we could tell.

But if the "Top-Level Switch → Disorientation" portion of this causal model is believable, we would expect that in the Local Context condition, where the top-level mechanism was also available for top-level switches, a relationship between top-level switches and disorientation would be found. However, neither frequency nor proportion of top-level switches predicted any of the three disorientation measures in the Local Context condition. The same is true of the Full Overview condition, where, in fact, the relationship between top-level switches and "I generally knew where I was in the Web site" approached a significantly *positive* relationship ($r = 0.587$, $p = 0.074$). Moreover, consider that distal movements, by the same presumption we made regarding top-level switches, are between source and destination pages with less related information topics than local transitions. Neither frequency nor proportion of distal transitions predicted any of the three disorientation measures, in both the Full Overview and the Local Context conditions. In short, all of this evidence points to one conclusion: moving amongst unrelated information topics in the Web space appears not to have been related to disorientation. However, that is based on the assumption that in fact the information topics were sufficiently unrelated within the stimulus site, which was not measured, and so further investigation is certainly needed. For the purposes of this discussion, we need only consider that top-level switches and distal movements in this study were not causes of disorientation.

We have good reason, then, to consider the two models above implausible. But there is another potentially problematic arrow: "Disorientation → Top-Level Switch," as part of the larger picture, "Disorientation → Top-Level Switch → Navigational Volatility." Note that this arrow is not inconsistent with navigational volatility causing disorientation. We could, rather, be faced with another potentially ugly cycle in Web navigation:

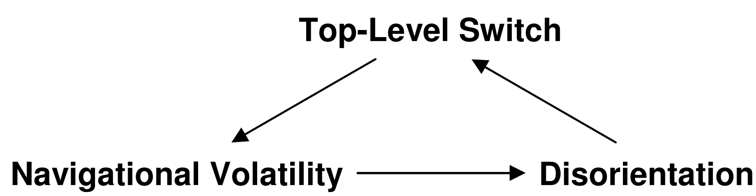

We might simply argue that disorientation causes navigational volatility directly, although this seems slightly less plausible. Notice that predictability now plays a role. In order for us to accept "Disorientation → Navigational Volatility" we would need to accept that disoriented users consciously recognize the hyperlinks that will lead to navigational changes, and decide that such change is desirable, leading to a transition. In this study, all hyperlinks within the same mechanism had homogenous volatility, and under such circumstances we might reasonably believe that a direct causal link from disorientation to navigational volatility is at least nearly as plausible as "Disorientation → Top-Level Switch → Navigational Volatility." In any event, if we are to believe that disorientation has any such causal power, we ought to be able to explain why it flexed it only in the Partial Overview condition and not the Local Context condition. (We could explain why it did not in the Full Overview condition, which did not include a top-level links mechanism, by stating simply that without such a homogenous mechanism, users were less able to make hyperlink decisions that would predictably move them to another sub-site, or at least a distal page.) There simply does not seem to be a good answer. If the top-level hyperlinks significantly "invite" disoriented users, we would expect the extent to which users invoke that particular mechanism to predict disorientation. Neither frequency nor proportion of top-level mechanism clicks predicted any of the three disorientation measures in the Local Context condition. (Note here that "I generally knew where I was in the Web site" *did* correlate negatively with both the frequency ($r = -0.681$, $p = 0.030$) and proportion ($r = -0.643$, $p = 0.045$) of home page hyperlink clicks. Thus, it was the home page link in the top-left corner of the page, just as we might expect, that was doing the "inviting" for disoriented users, not the top-level hyperlinks.)

Now the causal relationship appears more plausible: "Navigational Volatility → Disorientation." But we still need an explanation as to why navigational volatility flexed *its* causal power only in the Partial Overview condition and not in the Local Context condition. The simple and crucial observation is that the distributions of navigational volatility in the two conditions were quite different. In the Partial Overview condition, users encountered a sort of all or nothing change in their navigation options during the experimental session. In making local transitions, the Breadcrumb List mechanism would alter slightly for a small navigational volatility score, while in making distal movements between sub-sites, a left-column Partial Overview mechanism would present a whole new set of navigation options. In the Local Context condition, on the other hand, the changes were more subtle and spread out. Usage of the Sibling List mechanism resulted in a new set of child node hyperlinks, with distal movements resulting in lesser change than in the Partial Overview condition.

A reasonable conclusion seems to be that the all or nothing nature of users' navigational volatility distributions in the Partial Overview condition were more noticeable, and more disorienting, than the more subtle and graded changes typically encountered with a Local Context navigational scheme. And in fact we have a way of confirming the noticeable nature of the navigational changes: in the Local Context condition, actual navigational volatility did not correlate significantly with "The navigation options (links) seemed to stay the same as I navigated" ($r = -0.126$, $p = 0.728$), but *did* in the Partial Overview condition ($r = -0.648$, $p = 0.043$). Not surprising, since the changes were dramatic for Partial Overview participants when they did occur, but telling. A key point is: the navigational scheme used will affect the extent to which navigational volatility plays a role in disorientation, based upon the sorts of volatility distributions we can expect as a result of the scheme. The navigational changes, when dramatic enough, make a difference. Notice that habituation now has a role. The extent to which a user is habituated in a navigation patch may make the changes seem even more dramatic when they do occur, as we will discuss in section 17.5.

17.1.2 Site size and complexity

H2: A volatile navigation session will cause a user to perceive a site as being larger and more complex.

The study results suggest that navigational volatility does lead to increased site size perception, but do not suggest an effect on complexity perception. The result was significant in the Local Context condition, but not in the Partial Overview condition. One might expect that simply looking at the number of pages a navigator visited would tell us something about her perception of the site size, but in the Local Context condition this is not the case (although across all three conditions it is, with site size estimate correlating with number of unique pages visited ($r = 0.488$, $p = 0.006$)).

We mentioned in our disorientation discussion that navigational volatility appeared to go relatively unnoticed in the Local Context condition. However, there is a difference between (i) noticing changes in navigation options during the session, and (ii) as a result of those changes, being exposed to a greater variation of information topics in the site.

This brings us into the territory of perceived global coherence, the topic of H4. The basis of linking H2 (Navigational Volatility → Increased site size perception) and H4 (Navigational Volatility → Decreased perceived global coherence) was that a site which seemed smaller might, conceptually, seem to have a more "compact" set of information topics that are more tightly related, and vice versa. But this is not what happens. A significant effect was not found linking perceived global coherence of the site to its perceived size. The user's perception of the broad ranging nature of the site, which might be captured by the centrality ratings in Part V of the post-navigation questionnaire (Appendix C.5), also failed to predict site size perception.

This study suggests that, with Local Context navigation support, neither the exposure to more and more information topics in a site, nor the perception of the site as being broad-ranging in its set of topics, leads a user to view the site as larger, which, unfortunately, leaves this discussion somewhat at a loss. It may be the case that exposure to a broad range of the site's information topics is the causal factor (true of all three conditions, and of the Partial Overview and Local Context conditions combined), and that navigational volatility, even though relatively unnoticed with Local Context, leads to this exposure. However, further research will be necessary to understand these complex relationships.

17.1.3 Model of exploration

H3: A volatile navigation session will cause a user to have a less accurate model of her own exploration of a site.

This study did not confirm H3. Although within and across the Partial Overview and Local Context conditions navigational volatility correlated positively with the extent to which a user was in err in estimating her own exploration of the site's main sections, a significant effect was not found. Lateral movements and total unique page visits led to more accurate models, perhaps suggesting that breadth-first strategies have a positive effect.

In general, exploration models are emphasized here as an important area for future research. It is reasonable to believe that the behind the door problem of Web navigation remains one of the more crippling effects of poor navigation design. Users will undoubtedly benefit substantially from more accurate models of what paths they have followed and the extent to which they have followed them, navigation designers will benefit from knowing how to promote such understanding, and research will provide these answers.

17.1.4 Perceived global coherence

H4: A volatile navigation session will cause a user to perceive a site as being less globally coherent.

Interestingly, a volatile navigation session predicted a low level of perceived global coherence in the Partial Overview condition, but predicted a high level in the Local

Context condition. The explanation for the Partial Overview condition result is as was predicted: a highly volatile navigation session, when noticeable, leads a user to recognize connectivity differences amongst pages in the space, which leads her to believe the information topics of these pages are likely less related.

The results in the Local Context condition are trickier, and do not lend themselves to any obvious explanations. Distal pairs of Web pages in a site hierarchy, with relatively unrelated information topics, typically have relatively different navigational support — that is, relatively different connectivity. Yet the exposure to different navigational support was predictive of increased perceived global coherence; users tended to view the site topics as more tightly related anyway. As with site size perception, the number of unique pages visited did not predict increased perceived global coherence, so page visits provide no potential explanation. The most plausible explanation, it seems, is that exposure to the navigational differences (caused by navigational volatility) allowed users with Local Context support to see connections between distal pages they otherwise would not have seen, and so led them to view such pages as more related. In the Partial Overview condition, such volatility was not necessary to view many of the distal pages as connected, since the Partial Overview mechanism displayed many distal hyperlink pairs together as it was. However, even this explanation is highly speculative, and requires further research.

Note here that the extent to which a navigator views page titles as central to the main topic of the site may measure a user's mental model of the site structure more so than her perceived global coherence. Participants tended to rate page titles higher in the site hierarchy as more central, perhaps relying on their memory of where page topics appeared in the site. The relatedness ratings of site topics appears to be a more suitable metric, and is recommended for future studies assessing the concept.

17.2 Click-stream analysis

The click-stream data collected during the 15-minute experimental navigation session demonstrated many of the vast differences between the three navigational schemes investigated. This section will briefly put those results into some perspective.

17.2.1 Page visits

More navigational context, such as in the Full Overview condition, led to decreased page visits, consistent with previous work [Park+00a/b]. Similarly, the additional context led to decreased page revisitation, also consistent with previous work [DaniIP]. The mean revisitation rate of 53.28% is quite close to the rate reported by Tauscher and Greenberg [Taus+97a/b] in a global navigation study (58%), and the rate found in the control condition of a previous intra-site study (56.05%). As pointed out it that previous study, however, the revisitation rate is heavily influenced by Back button usage, and the close replication of the Tauscher and Greenberg rate is likely somewhat coincidental. Consistent with previous work [DaniIP], overview context caused navigators to dig deeper into the site hierarchy and spend more time there.

17.2.2 Method of access

Both the Partial Overview and Local Context schemes included a Top-Level navigational mechanism, and Partial Overview participants made more use of it. The reason is unclear, as the two conditions did not differ in their mean levels of perceived global coherence. One potential explanation, although only anecdotal, is that Partial Overview appeared to, in some cases, use the top-level hyperlinks as a way of scanning the terrain, clicking on a number of them consecutively, perhaps to get a sense of the scope of each sub-site. This could not be done as easily in the Local Context condition, since less of a sub-site's structure could be viewed from a top-level page, and thus digging into the structure would have been necessary for the same sub-site scanning Partial Overview participants were capable of. However, again, this explanation is speculative, and will benefit from further investigation.

Consistent with previous work [DaniIP], overview context resulted in decreased Back button usage. This result is closely related to the decreased revisitation rate resulting from the additional navigation support. Similar to global navigation behavior studies [Catl+95, Taus+97a/b], other browser-supported navigation, such as the Forward button and History List of the Back button, were rarely used. (Of the 1730 navigation acts across the 30 participants, the Forward button and History List were both used once.)

17.2.3 Relationship and direction

Overview context was generally used to make distal movements in the site hierarchy, and to dig deeper into the site structure. In previous work [Dani01, DaniIP], navigators provided with a constantly visible overview rarely made hierarchical path jumps down to a grandchild page or up to a grandparent page, but in this study such jumps were more common. There are three primary differences between the navigation support provided in the overview condition of that study, and the Full Overview condition in this study: (i) the previous study provided Focus+Context support, while the control of this study provided only a Full Overview without local context navigation aids, (ii) the previous study's constantly visible overview did not include a current location marker, and (iii) the previous study's overview was presented in a separate frame. Any of these differences could account for the increase in hierarchical path jumps. First, the lack of context support in this study's control may have caused navigators to focus on the overview for local movement more than they otherwise would, and since it includes a hyperlink to the current page's grandchild, it might lead to more path jumps than focusing on local context and relying on the overview only when a distal movement is desirable would have. Second, the current location marker, which was the proposed explanation for the lack of path jumps in the previous study, may have allowed the user to be more aware of her current location, and in particular of what path she was currently headed down. Third, the use of frames keeps the local context of the current page in focus above the scroll line, while without the separate frame the viewable hyperlinks after a transition are those with the highest list position (resulting in view volatility).

17.2.4 Duration of page visit

Although no claims can be made regarding online reading and scanning based solely on click-stream data, the word count of a Web page had a clear impact on the duration of page stay, with longer word count pages resulting in more time spent. Not surprisingly, users tended to spend less time on a page when visited via the Back button, likely due to rapid, successive backtracking and "hub and spoke" behavior [Catl+95].

Recall that the experimental navigational schemes used in this study were presented such that transitional attributes other than navigational volatility could be accounted for. Likely as a result, distal movements and top-level switches did not result in more time being spent at the destination page, unlike previous work in which such movements included other visual changes such as new color coding and different content presentation; this study suggests that navigation option changes that occur as a result of a page-to-page transition do not cause longer page stays. A possible explanation is that page stay depends more upon initial visual changes, and we know that users tend to look to the page content when they first arrive at a destination page.

17.2.5 List position bias

As with online reading behavior, no claims can be made regarding a navigator's willingness to scroll based solely upon click-stream data, but there was a clear list position bias in the Full Overview condition, with participants in that condition being more inclined to invoke hyperlinks high up in the overview list, above the scroll line. The mean position was greater in the Full Overview condition than in the others. Note here that one cannot look simply to the extent to which navigators spend time above the scroll line or invoke links above the scroll line, without a comparison to another condition in which the information need was the same. Doing so would miss a crucial point: a bias might only point to the relative importance of navigation options above the scroll line, and so would not point to any bias caused by the navigational scheme. Because navigators in the Partial Overview and Local Context conditions had the same tasks, we can reasonably conclude that the tendency for navigators with a Full Overview (not in a separate frame) to invoke links with high list positions resulted from the differences between the available mechanisms. We note that one of the most salient differences was the fact that the Full Overview extended well below the scroll line, while no navigation options in the other two conditions appeared below it.

17.3 Usability and performance

In each condition, the left-column mechanisms tended to be rated as subjectively more useful; it is unknown whether this speaks to some fundamental preference for left-column mechanisms over, for example, horizontal mechanisms across the top of the page.

Greater context did appear to lead to decreased task abandonment, consistent with previous work [DaniIP], although task abandonment in the Full Overview did not differ significantly from that in Local Context condition. The difference between the navigational schemes of the two studies most likely responsible is the lack of a Focus+Context approach in the control of this study; the positive effects of Focus+Context presentation were briefly discussed in Chapter 4.

17.4 Mental models

This study suggests that there is more to site size perception than the amount of navigational support provided. There undoubtedly remain other tricks site designers employ to make a small site look larger and more impressive, but in this study such tactics were accounted for. The navigational scheme of the Partial Overview condition was such that the (nearly) maximal number of hyperlinks were presented along the left side of the screen at all times, without any scrolling requirement for navigation. This characteristic may have contributed to Partial Overview participants estimating the site as significantly larger than those with other navigational schemes. Strongly hierarchical sites are viewed as smaller than weakly hierarchical sites [Modj+97], and this study, moreover, suggests that strongly hierarchical sites may be viewed as larger when the visibility of available navigation options is maximized.

17.5 Transitional effects

Most of the transitional effects under consideration were discussed in relation to the experimental hypotheses, in section 17.1. Interestingly, navigational volatility also appeared to lead to decreased usability in the Partial Overview condition, where, as we have mentioned, the navigation changes were noticeable.

Recall from Chapter 13 the suggestion that the user's experience with the Web site should be the cause of any discrepancy between actual and perceived volatility. Specifically, habituation should lead to overestimation of actual volatility, and predictability should lead to underestimation of actual volatility. The metrics for navigational volatility, habituation, and predictability did exactly what we would expect them to do in the Local Context Overview condition, but only volatility and habituation did as expected in the Partial Overview condition, not predictability. That is, with partial overview support, navigational predictability did not lead to an underestimation of actual volatility. A possible explanation is that, given the higher level of predictability of a navigational scheme with a persistent overview, it is not clear a designer could cause users to further underestimate the navigational volatility of their experience with the site. The extent to which users do in fact become habituated within navigation patches and have predictive power at the source pages of transitions remains an open area for research.

17.6 Applicability

The applicability of a Web navigation study can be thought about along four dimensions: user, task, site, and design scheme.

User. In this study, participants likely had generally high Web expertise, were not domain experts, and were likely more efficient information seekers than a wider age sample would be. Moreover, the results may be restricted to Western culture, given the bases for the navigational schemes investigated [vanS+01].

Task. This study investigated directed search tasks with low complexity. The navigation session imposed some time pressure.

Site. This study applies to relatively small sites, or sub-sites of larger ones, which are hierarchically organized. The study results may not apply to sites with very high levels of internal or external structural consistency.

Design scheme. The design schemes investigated in this study were common to hierarchically organized sites, had generally homogenous screen real estate usage, and were not subject to slow download times. Note that design scheme decisions were generally "by the book," based upon previous research in the field. Sites that deviate heavily from the basic, tested design principles will not necessarily be informed by these results.

These specifications are indicative of a fairly common circumstance on the Web: a general Web expert needs specific information on a topic she is generally unfamiliar with, and navigates within a small hierarchical site or sub-site.

There is of course a fifth, implicit applicability dimension to be considered, namely era. Web navigation studies can often produce results that are to some extent artifacts of the time period. Back button usage provides a nice case study. Catledge and Pitkow [Catl+95] reported that Back button clicks accounted for 41% of all navigation acts, based on data collected in 1994. Tauscher and Greenberg [Taus+97a/b] later reported the number at 30% (and the revisitation rate found in their work is arguably an artifact of the era of their study [Cock+01, McKe+01]). The cause of this apparent era dependency might be found in any of the four dimensions above. The general Web expertise of users might be changing such that backtracking becomes less common, the sorts of tasks people complete on the Web might be changing so as to cause less backtracking, sites' organizations and structures might be changing in such a way as to have the effect, or widespread changes in the design schemes typically employed may be the cause. Notice that if design schemes evolved from the time of the Catledge and Pitkow [Catl+95] work to the time of the Tauscher and Greenberg [Taus+97a/b] work to include more extensive overview support, empirical evidence suggests this shift would in fact lead to decreased Back button use.

Any of the specifications of the four dimensions above could be more or less important depending on the time period. For example, if hierarchical organizations became uncommon, this study would be subject to such era dependency. An effort was made in this work to avoid any such dependency, focusing on a scenario that is likely to remain common. For example, hierarchies tend to match the way we think about the world and navigational spaces [Stev+78, Chas83].

18 Conclusions

Web navigation research is particularly challenging not only because the Web is a massive, rapidly changing information ecology, but because the complexities of the navigational constructs, available information, design schemes, potential tasks, users, and their many relationships are utterly daunting. Supporting users and their tasks — and supporting Web information designers in doing so — will require, to use the navigational metaphor, a map of the research space and its many considerations. The study and discussion in this thesis hopefully make a useful contribution to that space.

18.1 Considerations

Web information designers and site architects benefit substantially from having a battery of usability and design considerations, and empirical data to reveal the nature and significance of those considerations. This thesis provides empirical data regarding the impact of navigational volatility, and suggests that it has effects worth noting.

Of course, there is, quite simply, no Holy Grail consideration of disorientation and mental models on the Web. Navigational volatility is undoubtedly one of many contributing factors. Web considerations ought to progress from general intuitions to empirically validated concerns, and finally to considerations that can be efficiently and effectively incorporated into the Web development and evaluation process. This thesis is believed to be a small step in that process for a particular kind of consideration, and hopefully may serve information designers well.

18.2 Limitations

The applicability discussion in the previous chapter provided a broad sense of the limitations of this work. As will often be necessary in Web navigation research, a small piece of the huge and complex four-dimensional space described there was investigated in this study, leaving open, in particular, a wide range of task types. Similarly, as mentioned above, this study looked only at a few pieces of a larger picture of transitional disorientation, proposed in Chapter 13.

Another limitation also comes out of a choice of focus in this study; we could not discuss performance measures since participants were given the option to abandon a task, with some users possibly being more inclined to abandon difficult tasks; thus, the simple measure of completed tasks could not be reliably analyzed. This decision was based on the already large body of performance research, and need for more investigations of other kinds of behavioral effects, but nonetheless limits our discussion to other usability concerns.

18.3 Future research

A number of potential investigations immediately present themselves. First and foremost, the applicability limitations of the work could be addressed with studies considering, for example, general purpose browsing tasks. Similarly, other pieces of the disorientation model might be investigated as well, in particular the effects of content volatility, likely more immediately noticeable, and closely related to the notions of visual turbulence and visual momentum. The effects of these concepts on the Web deserve investigation.

We have noted a number of areas also deserving of further research throughout our discussion. The volatility associated with some types of dynamic hyperlinks (especially those based upon navigational history), for example, deserves attention. The user's exploration model will also make for interesting future research, as well as how navigation designers might better support forward navigation and predictability.

Finally, we might consider how the considerations put forth in this work might eventually be useful in the Web evaluation process. Given likely user paths through an information space, usability evaluations might benefit from considering the navigational volatility of those paths, and the specific effects described in this work.

18.4 Summary

We have mapped out in our discussion a space of attributes and relationships involved in the Web navigation process, and discussed a few key ways in which the extremely complex interaction between user, task, and Web site might be measured and assessed. Finally, we discussed an empirical study of transitional volatility, primary concerned with (i) the navigational and content changes of the Web interface in page-to-page transitions, and (ii) the user's ability to reorient herself to these changes, focusing on the navigational changes in particular.

The extent to which a user's session is navigationally volatile affects her level of disorientation, the usability of the site, and the way in which the site is subjectively viewed, interacting with the navigation support provided.

APPENDICES, REFERENCES, AND ACKNOWLEDGMENTS

Background and Instructions

You have a number of legal questions you would like answers to, and have decided to make use of a "self-help" Web site containing practical information about legal matters. You have 15 minutes of spare time during which you will navigate the Web site trying to find answers to your legal questions as best you can.

You will be given a stack of cards, face down, each of which contains a legal question; the answers to these questions are to be found on the Web site you will navigate. Attempt to answer only one question at a time by turning the card over.

If you believe you have discovered the answer, highlight the information you believe answers the question, as though you were to copy and paste the information into a document to read later (but you should not actually copy and paste it; the screen output will be recorded, so it will not be necessary to do so).

After you highlight the answer, click on the "ANSWER QUESTION" link in the buttons bar of the browser, and follow the link provided on that page back to the home page of the Web site. This will "stamp" your completion of the question. Then turn the next card over to begin the next question.

Make an effort to answer the questions as best you can, but if you become discouraged and do not believe you will find the answer, you may move on to the next question. Do this by clicking on the "SKIP QUESTION" link in the buttons bar of the browser, and continuing as you would if you had answered the question. Once you skip a question, you cannot return to it later.

Navigate the site for 15 minutes, answering the questions as best you can, at which time I will alert you to stop.

Questionnaire Instructions

Now that you have navigated the Web site, you will be asked to provide a series of ratings and brief written responses to questions regarding the experience. This section of the study will likely take about 20 minutes, but there is no time pressure. Simply provide the requested ratings and brief answers as honestly and accurately as you can. The questionnaire is broken up into five short parts, which you will be asked to complete in order. Each time you complete a part, alert me and I will provide you with the next one.

Navigation Questionnaire PART I

Please estimate the size of the site you just navigated in number of pages: _______

Of the number of pages you estimated above, estimate how many of those pages you visited: _______

As mentioned on the home page, the site had five main sections: Going to Court, Family Law, Domestic Violence, Juvenile Law, and Small Claims. Please estimate how much of each of those sections you explored. In each space, provide an estimate between 0 and 100%:

Going to Court:	______%	Family Law:	______%
Domestic Violence:	______%	Juvenile Law:	______%
Small Claims:	______%		

Please indicate how often (on a 1-10 scale) you believe the following occurred while you navigated, with a rating of 1 meaning that you believe it never happened, and a rating of 10 meaning that it seemed to happen very frequently:

I desired to return to a previously visited page, and successfully returned to it:

← Not Often Very Often →

1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10

I desired to return to a previously visited page, and was unable to find it:

← Not Often Very Often →

1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10

Briefly provide any specifics you recall of the instances in which you desired to return to a previously visited page, such as what steps you took to try to find the page:

Please indicate how often (on a 1-10 scale) you believe the following occurred while you navigated, with a rating of 1 meaning that you believe it never happened, and a rating of 10 meaning that it seemed to happen very frequently:

I clicked on a link based on its position on the screen, without needing to read the link title:

← Not Often Very Often →

1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10

I clicked on a link based on its position on the screen, without reading the link title, and discovered the link took me somewhere other than where I expected:

← Not Often Very Often →

1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10

Briefly provide any specifics you recall of the instances in which you clicked on a link based on screen location, such as:

The names of the pages the links took you to; and/or

The areas of the screen (i.e. “bottom-left”) where the links appeared

Navigation Questionnaire PART II

Please indicate how much you agree or disagree (on a 1-10 scale) with the statements made below, by checking the box for your rating:

Your Ratings

	← Strongly Disagree									Strongly Agree →
1. I generally knew where I was in the Web site	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
2. The information in the site was well organized	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
3. The navigation options (links) were well organized on the page	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
4. The information in the site was credible and believable	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
5. It was easy to find the information I needed in the site	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
6. The information in the site seemed to be tied together and well connected	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
7. The (underlined) link titles well summarized the information they led to	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
8. I needed to keep track of a lot of things in order to find the information	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
9. The number of navigation options (links) was overwhelming	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
10. It was easy to predict where in the site clicking on a link would take me	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
11. I felt lost and needed to reorient myself	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
12. I felt trapped and unable to easily move to unexplored areas of the site	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
13. It was easy to understand how the page would change when I clicked on a link	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
14. When I felt lost, it was easy to reorient myself in the site	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
15. The site seemed large	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
16. The site seemed complex	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
17. The site was easy to navigate	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
18. The navigation options (links) seemed to stay the same as I navigated	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
19. It was easy to return to pages I had previously visited	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
20. The topics in the site were related to one another and pertained to one coherent topic	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10

Navigation Questionnaire PART III

In the site you navigated, each page you visited had a number of navigational options (links). The options you had were in one of three areas on the page:

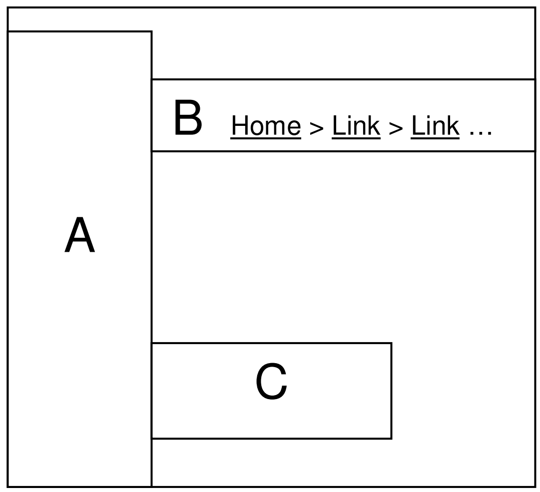

Section B had links that looked something like "Home > Link > Link." For sections A, B, and C, please briefly describe as best you can the purpose of that section, including how you believe the links in that section were related to one another and to the page you were at:

Section A:

Section B:

Section C:

In the top-left corner of the page (just above section A), there was a link called "Self-Help Center Home." Please rate how useful (on a 1-10 scale) you found the sections in the above figure to be in looking for the information, and how useful you found the "Self-Help Center Home" link to be, by checking the box for your rating:

		← Not Useful					Very Useful →			
Section A:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section B:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section C:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
"Self-Help Center Home" Link:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10

Finally, please rate how useful you found all of the navigation options together to be in looking for the information:

		← Not Useful					Very Useful →			
All of my navigation options:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10

Navigation Questionnaire PART III

In the site you navigated, each page you visited had a number of navigational options (links). The options you had were in one of four areas on the page:

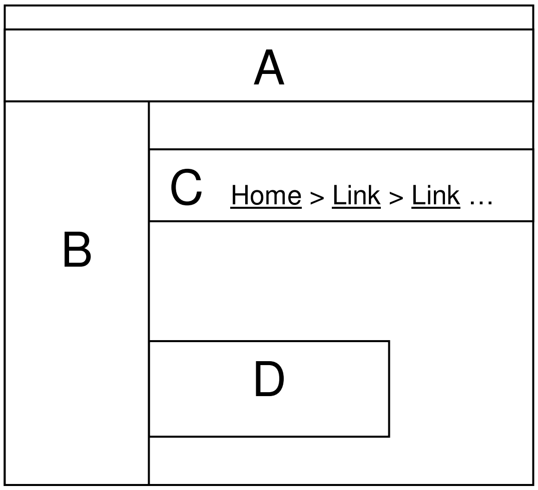

Section C had links that looked something like "Home > Link > Link." For sections A, B, C, and D, please briefly describe as best you can the purpose of that section, including how you believe the links in that section were related to one another and to the page you were at:

Section A:

Section B:

Section C:

Section D:

In the top-left corner of the page (just above section A), there was a link called "Self-Help Center Home." Please rate how useful (on a 1-10 scale) you found the sections in the above figure to be in looking for the information, and how useful you found the "Self-Help Center Home" link to be, by checking the box for your rating:

		← Not Useful					Very Useful →		
Section A:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section B:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section C:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section D:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
"Self-Help Center Home" Link:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐ 10

Finally, please rate how useful you found all of the navigation options together to be in looking for the information:

		← Not Useful					Very Useful →		
All of my navigation options:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐ 10

Navigation Questionnaire PART III

In the site you navigated, each page you visited had a number of navigational options (links). The options you had were in one of five areas on the page:

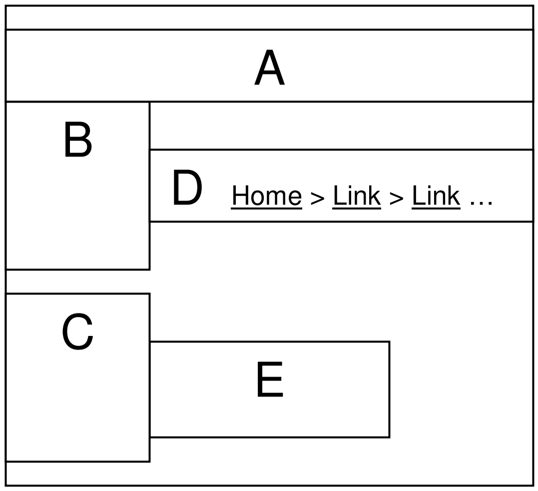

Section D had links that looked something like "Home > Link > Link." For sections A, B, and C, please briefly describe as best you can the purpose of that section, including how you believe the links in that section were related to one another and to the page you were at:

Section A:

Section B:

Section C:

Section D:

Section E:

In the top-left corner of the page (just above section A), there was a link called "Self-Help Center Home." Please rate how useful (on a 1-10 scale) you found the sections in the above figure to be in looking for the information, and how useful you found the "Self-Help Center Home" link to be, by checking the box for your rating:

		← Not Useful					Very Useful →			
Section A:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section B:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section C:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section D:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
Section E:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10
"Self-Help Center Home" Link:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10

Finally, please rate how useful you found all of the navigation options together to be in looking for the information:

		← Not Useful					Very Useful →			
All of my navigation options:	1 ❐	❐	❐	❐	❐	❐	❐	❐	❐	❐ 10

Navigation Questionnaire PART IV

Please indicate how related or unrelated to one another (on a 1-10 scale) you believe the following pairs of topics from the site you navigated are, by checking the box for your rating:

Topic 1	*Topic 2*	Your Ratings: ← Unrelated … Related →
Family Law	Small Claims	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Appeal, collect, or vacate judgment	Prepare, file, and serve a claim	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Court clerk's office	Domestic violence clinics	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Juvenile Law	Special education	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Court hearing preparation	Appellate courts	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Custody and visitation	Juvenile court processes	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Court fees and fee waivers	Statute of limitations	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Mediation	Child and spousal support	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Types of restraining orders	Probation department	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Arbitration	How to change a court date	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Being tried as an adult	Delinquency	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Divorce, separation, annulment	Voluntary declaration of paternity	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Who must appear in court	How to countersue	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
District attorney's office	Domestic Violence	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10
Small Claims	Collecting a judgment	1 ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ ❒ 10

Navigation Questionnaire PART V

Please indicate how central or peripheral to the main topic of the site (on a 1-10 scale) you believe the following topics are, by checking the box for your rating.

	Your Ratings									
	← Peripheral								Central →	
Family Law	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
Change a hearing date or claim	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
Probation department	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
Small Claims	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
At the courthouse	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
Being tried as an adult	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
Going to Court	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
Restraining orders	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10
Focusing on the child's needs	1 ❒	❒	❒	❒	❒	❒	❒	❒	❒	❒ 10

Task Orders

ORDER 1

Where can you get information on local department locations and court hours?
For how long can a restraining order last?
You want to change the date of a hearing but don't have time to mail a request; what should you do?
Who can provide assistance with small claims procedures and law?
What court hears child custody cases?
Who makes the decision to press or drop charges in a criminal case?
What kind of evidence should you bring to a court hearing?
What kinds of offenses can typically lead to a child being tried as an adult?
What are the official reasons for which you can legally end a marriage?
What can a family law facilitator do for you?
Under what circumstances can you get a child back who has been taken by a social worker?
Someone you sued and beat in court is supposed to pay you, but won't; what can you do?
What will a parole agent do if you commit a parole violation?
Under what circumstances can a child support order be changed?
You have technical questions about a case; what is your best court option?
Under what circumstances can you call the police to keep the peace?
Who should you consult for information about financial obligations between unmarried couples?
How many of the witnesses from a case's first hearing must show up for an appeal?
What information will a presentence investigation report for the court include?
Who can you contact to find out if there is a domestic violence program in your county?
What often determines how much it will cost to hire a process server?
What should be included in any good parenting plan document for the court?
What is the filing fee for a small claims case?
What must someone do while under a residence exclusion order?
What would entitle a child to special education services?

ORDER 2

You have technical questions about a case; what is your best court option?
What are the official reasons for which you can legally end a marriage?
What would entitle a child to special education services?
What is the filing fee for a small claims case?
Who should you consult for information about financial obligations between unmarried couples?
Who makes the decision to press or drop charges in a criminal case?
What court hears child custody cases?
What must someone do while under a residence exclusion order?
What kinds of offenses can typically lead to a child being tried as an adult?
What should be included in any good parenting plan document for the court?
Where can you get information on local department locations and court hours?
You want to change the date of a hearing but don't have time to mail a request; what should you do?
What will a parole agent do if you commit a parole violation?
Under what circumstances can you get a child back who has been taken by a social worker?
For how long can a restraining order last?
Under what circumstances can a child support order be changed?
What information will a presentence investigation report for the court include?

What often determines how much it will cost to hire a process server?
Who can you contact to find out if there is a domestic violence program in your county?
What kind of evidence should you bring to a court hearing?
Under what circumstances can you call the police to keep the peace?
Someone you sued and beat in court is supposed to pay you, but won’t; what can you do?
What can a family law facilitator do for you?
How many of the witnesses from a case’s first hearing must show up for an appeal?
Who can provide assistance with small claims procedures and law?

References

[Aalt+98] Aaltonen, A., A. Hyrskykari, & K. Räihä (1998). 101 spots, or how do users read menus? *CHI '98, Human Factors in Computing Systems*, 132-139.

[Ahuj+01] Ahuja, J.S. & J. Webster (2001). Perceived disorientation: an examination of a new measure to assess Web design effectiveness. *Interacting with Computers, 14(1)*, 15-29.

[Albe97] Albers, M.J. (1997). Cognitive strain as a factor in effective document design. *ACM Systems Documentation '97*, 1-6.

[Alle94] Allen, B. (1994). Perceptual speed, learning, and information retrieval performance. *ACM Research and Development in Information Retrieval '94*, 71-80.

[Ande80] Anderson, J.R. (1980). *Cognitive Psychology and its Implications*. San Francisco, CA: W.H. Freeman.

[Babi+91] Babiker, E.M., H. Fujihara, & C.D.B. Boyle (1991). A metric for hypertext usability. *ACM Systems Documentation '91*, 95-104.

[Bach+97] Bachiochi, D., M. Berstene, E. Chouinard, N. Conlan, M. Danchak, T. Furey, C. Neligon, & D. Way (1997). Usability studies and designing navigational aids for the World Wide Web. *Sixth International World Wide Web Conference*.

[Bart+95] Bartram, L., A. Ho, J. Dill, & F. Henigman (1995). The continuous zoom: a constrained fisheye technique for viewing and navigating large information spaces. *ACM User Interface Software and Technology '95*, 207-215.

[Beny+97] Benyon, D. & K. Höök (1997). Navigation in information spaces: supporting the individual. *IFIP INTERACT '97*, 39-46.

[BernM98] Bernstein, M. (1998). Patterns of hypertext. *ACM Hypertext '98*, 21-29.

[BernM00] Bernstein, M. (2000). More than legible: on links that readers don't want to follow. *ACM Hypertext '00*, 216-217.

[BernT+94] Berners-Lee, T., R. Cailliau, A. Luotonen, H.F. Nielsen, & A. Secret (1994). The World-Wide Web. *Communications of the ACM, 37(8)*, 76-82.

[Bieb+94] Bieber, M. & J. Wan (1994). Backtracking in a multiple-window hypertext environment. *ACM European Conference on Hypermedia Technology '94*, 158-166.

[Björ+99] Björk, S., L.E. Holmquist, & J. Redström (1999). A framework for Focus+Context visualization. *IEEE Information Visualization '99*, 53-57.

[Bota+91] Botafogo, R.A. & B. Shneiderman (1991). Identifying aggregates in hypertext structures. *ACM Hypertext '91*, 63-74.

[Bota+92] Botafogo, R.A., E. Rivlin, & B. Shneiderman (1992). Structural analysis of hypertexts: identifying hierarchies and useful metrics. *ACM Transactions on Information Systems, 10(2)*, 142-180.

[Bray96] Bray, T. (1996). Measuring the Web. *Fifth International World Wide Web Conference*.

[Broo98] Brooks, T.A. (1998). The Semantic Distance Model of relevance assessment. *Proceedings of the ASIS Annual Meeting, 35*, 33-44.

[Brow88] Brown, P.J. (1988). Linking and searching within hypertext. *Electronic Publishing, 1(1)*, 45-53.

[Brow89] Brown, P.J. (1989). Do we need maps to navigate round hypertext documents? *Electronic Publishing, 2(2)*, 91-100.

[Brun+01] Brunk, B. & G. Marchionini (2001). Toward an Agileviews WWW sitemap kit: the generalized relation browser. Technical Report TR-2001-06, School of Information and Library Science, University of North Carolina, Chapel Hill.

[Bush45] Bush, V. (1945). As we may think. *Atlantic Monthly, 176*, 101-108.

[Byrn+99a] Byrne, M.D., J.R. Anderson, S. Douglas, & M. Matessa (1999). Eye tracking the visual search of click-down menus. *CHI '99, Human Factors in Computing Systems*, 402-409.

[Byrn+99b] Byrne, M.D., B.E. John, N.S. Wehrle, & D.C. Crow (1999). The tangled Web we wove: a taskonomy of WWW use. *CHI '99, Human Factors in Computing Systems*, 544-551.

[Calv97] Calvi, L. (1997). Navigation and disorientation: a case study. *Journal of Educational Multimedia and Hypermedia, 6*, 305-320.

[CampF+89] Campagnoni, F.R. & K. Ehrlich (1989). Information retrieval using a hypertext-based help system. *ACM Transactions on Information Systems, 7(3)*, 271-291.

[CampC+99] Campbell, C.S. & P.P. Maglio (1999). Facilitating navigation in information spaces: Road-signs on the World Wide Web. *International Journal of Human-Computer Studies, 50*, 309-327.

[Cant+85] Canter, D., R. Rivers, & G. Storrs (1985). Characterizing user navigation through complex data structures. *Behaviour and Information Technology, 4(2)*, 93-102.

[Card84] Card, S.K. (1984). Visual search of computer command menus. In H. Bouma & D.G. Bouwhuis (Eds.), *Attention and Performance X: Control of Language Processes*. London: Lawrence Erlbaum Associates, 97-108.

[Card+80] Card, S.K. T.P. Moran, & A. Newell (1980). The keystroke-level model for user performance time with interactive systems. *Communications of the ACM, 23(7)*, 396-410.

[Card+83] Card, S.K., T.P. Moran, & A. Newell (1983). *The Psychology of Human-Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[Card+86] Card, S.K., T.P. Moran, & A. Newell (1986). The model human processor. In K.R. Boff, L. Kaufman, & J.P. Thomas (Eds.), *Handbook of Perception and Human Performance, Vol. II*. New York: John Wiley & Sons, chapter 45.

[Card+96] Card, S.K., G.G. Robertson, & W. York (1996). The WebBook and the WebForager: an information workspace for the World-Wide Web. *CHI '96, Human Factors in Computing Systems*, 111-117.

[Card+99] Card, S.K., J.D. Mackinlay, & B. Shneiderman (Eds.) (1999). *Readings in Information Visualization: Using Vision to Think*. San Francisco, CA: Morgan Kaufmann.

[Card+01] Card, S.K., P. Pirolli, M. Van Der Wege, J.B. Morrison, R.W. Reeder, P.K. Schraedley, & J. Boshart (2001). Information scent as a driver of Web behavior graphs: results of a protocol analysis method for Web usability. *CHI '01, Human Factors in Computing Systems*, 498-505.

[Care+90] Carey, T., W. Hunt, & A. Lopez-Suarez (1990). Roles for tables of contents as hypertext overviews. *IFIP INTERACT '90*, 581-586.

[Carm+92] Carmel, E., S. Crawford, & H. Chen (1992). Browsing in hypertext: a cognitive study. *IEEE Transactions on Systems, Man, and Cybernetics, 22*, 865-883.

[Carr98] Carroll, J.M. (1998). Reconstructing Minimalism. In J.M. Carroll (Ed.), *Minimalism Beyond the Nurnberg Funnel*. Cambridge, MA: MIT, 1-18.

[Carr+87] Carroll, J.M. & M.B. Rosson (1987). Paradox of the active user. In J.M. Carroll (Ed.), *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction*. Cambridge, MA: MIT Press, 80-111.

[Catl+95] Catledge, L. & J. Pitkow (1995). Characterizing browsing strategies in the World Wide Web. *Computer Networks and ISDN Systems, 27*, 1065-1073.

[Char87] Charney, D. (1987). Comprehending non-linear text: the role of discourse cues and reading strategies. *ACM Hypertext '87*, 109-120.

[Chas83] Chase, W.G. (1983). Spatial representations of taxi drivers. In D.R. Rogers & J.A. Sloboda (Eds.), *Acquisition of Symbolic Skills*. New York: Plenum, 391-405.

[Chen97] Chen, C. (1997). Structuring and visualizing the WWW by generalized similarity analysis. *ACM Hypertext '97*, 177-186.

[Chen00] Chen, C. (2000). Individual differences in a spatial-semantic virtual environment. *Journal of the American Society for Information Science, 51(6)*, 529-542.

[Chen+96] Chen, C. & R. Rada (1996). Interacting with hypertext: a meta-analysis of experimental studies. *Human-Computer Interaction, 11(2)*, 125-156.

[Chi+00] Chi, E., P. Pirolli, & J. Pitkow (2000). The scent of a site: a system for analyzing and predicting information scent, usage, and usability of a Web site. *CHI '00, Human Factors in Computing Systems*, 400-407.

[Chi+01] Chi, E., P. Pirolli, K. Chen, & J. Pitkow (2001). Using information scent to model user information needs and actions on the Web. *CHI '01, Human Factors in Computing Systems*, 490-497.

[Chim+94] Chimera, R. & B. Shneiderman (1994). An exploratory evaluation of three interfaces for browsing large hierarchical tables of contents. *ACM Transactions on Information Systems, 12(4)*, 383-406.

[Chin+88] Chin, J.P., V.A. Diehl, & K.L. Norman (1988). Development of an instrument measuring user satisfaction of the human-computer interface. *CHI '88, Human Factors in Computing Systems*, 213-218.

[Choo+00] Choo, C.W., B. Detlor, & D. Turnbull (2000). Information seeking on the Web: an integrated model of browsing and searching. *First Monday, 5(2)*.

[Clar96] Clark, H.H. (1996). *Using Language*. Cambridge, England: Cambridge.

[Cock+96] Cockburn, A. & S. Jones (1996). Which way now? Analysing and easing inadequacies in WWW navigation. *International Journal of Human-Computer Studies, 45(1)*, 105-129.

[Cock+97] Cockburn, A. & S. Jones (1997). Design issues for World Wide Web navigation visualisation tools. *Proceedings of RIAO'97: The Fifth Conference on Computer-Assisted Research of Information*, 55-74.

[Cock+99] Cockburn, A. & S. Greenberg (1999). Issues of page representation and organization in Web browser's revisitation tools. *Proceedings of the OZCHI '99 Australian Conference on Human Computer Interaction*.

[Cock+01] Cockburn, A. & B. McKenzie (2001). What do Web users do? An empirical analysis of Web use. *International Journal of Human-Computer Studies, 54(6)*, 903-922.

[Comb+96] Comber, T. & J.R. Maltby (1996). Evaluating usability of screen designs with layout complexity. *Proceedings of the Third International Eurographics Workshop on Design, Specification, and Verification of Interactive Systems*.

[Conk87] Conklin, J. (1987). Hypertext: an introduction and survey. *IEEE Computer, 20(7)*, 17-41.

[Cove+88] Cove, J.F. & B.C. Walsh (1988). Online text retrieval via browsing. *Information Processing and Management, 24*, 31-37.

[Dahl+96] Dahlbäck, N., K. Höök, & M. Sjölinder (1996). Spatial cognition in the mind and in the world: the case of hypermedia navigation. *The 18th Annual Meeting of the Cognitive Science Society*, University of California, San Diego.

[DaniIP] Danielson, D.R. (in press). Web navigation and the behavioral effects of constantly visible site maps. *Interacting with Computers*.

[Dani01] Danielson, D.R. (2001). Web navigation and the behavioral effects of constantly visible site maps. Undergraduate honors thesis, Symbolic Systems Program, Stanford University.

[Dark+93] Darken, R.P. & J.L. Sibert (1993). A toolset for navigation in virtual environments. *ACM User Interface Software and Technology '93*, 157-165.

[Dee-L+95] Dee-Lucas, D. & J.H. Larkin (1995). Learning from electronic texts: effects of interactive overviews of information access. *Cognition and Instruction, 13(3)*, 431-468.

[DeRo89] DeRose, S.J. (1989). Expanding the notion of links. *ACM Hypertext '89*, 249-257.

[DeYo90] DeYoung, L. (1990). Linking considered harmful: designing and reading hyperdocuments. *Proceedings of the First European Conferences on Hypertext*, 238-249.

[Diap+00] Diaper, D. & P. Waelend (2000). World Wide Web working whilst ignoring graphics: good news for Web page designers. *Interacting with Computers, 13(2)*, 163-181.

[Dias+97] Dias, P & P. Sousa (1997). Understanding navigation and disorientation in hypermedia learning environments. *Journal of Educational Multimedia and Hypermedia, 6*, 173-185.

[Dias+99] Dias, P., M.J. Gomes, & A.P. Correia (1999). Disorientation in hypermedia environments: mechanisms to support navigation. *Journal of Educational Computing Research, 20(2)*, 93-117.

[Dieb96] Dieberger, A. (1996). Browsing the WWW by interacting with a textual virtual environment: a framework for experimenting with navigational metaphors. *ACM Hypertext '96*, 170-179.

[Dijk68] Dijkstra, E.W. (1968). Go to statement considered harmful. *Communications of the ACM, 11(3)*, 147-148.

[Dill+90] Dillon, A., C. McKnight, & J. Richardson (1990). Navigation in hypertext: a critical review of the concept. *IFIP INTERACT '90*, 587-592.

[Dufr+97] Dufresne, A. & S. Turcotte. (1997). Cognitive style and its implications for navigation strategies. *Artificial Intelligence in Education: Knowledge and Media in Learning Systems*, 287-293.

[Edwa+89] Edwards, D.M. & L. Hardman (1989). Lost in hyperspace: cognitive mapping and navigation in a hypertext environment. In R. McAleese (Ed.), *Hypertext: Theory and Practice*. Oxford, England: Intellect Books, 105-125.

[Egid+88] Egido, C. & J. Patterson (1988). Pictures and category labels as navigational aids for catalog browsing. *CHI '88, Human Factors in Computing Systems*, 127-132.

[Eick+93] Eick, S.G. & G.J. Wills (1993). Navigating large networks with hierarchies. *IEEE Visualization '93*, 204-209.

[Elm+85] Elm, W.C. & D.D. Woods (1985). Getting lost: a case study in interface design. *Proceedings of the Human Factors Society 29th Annual Meeting*, 927-931.

[Enge62] Engelbart, D.C. (1962). Augmenting human intellect: a conceptual framework. Summary Report AFOSR-3223 under Contract AF 49(638)-1024, SRI Project 3578 for Air Force Office of Scientific Research, Stanford Research Institute.

[Enge+68] Engelbart, D.C. & W.K. English (1968). A research center for augmenting human intellect. *AFIPS Conference Proceedings, 33*, 395-410.

[Ersk+97] Erskine, L.E., D.R.N. Carter-Tod, & J.K. Burton (1997). Dialogical techniques for the design of Web sites. *International Journal of Human-Computer Studies, 47*, 169-195.

[Farr+01] Farris, J.S., K.S. Jones, & P.D. Elgin (2001). Mental representations of hypermedia: an evaluation of the spatial assumption. *Proceedings of the Human Factors and Ergonomics Society 45th Annual Meeting*.

[Fein88] Feiner, S. (1988). Seeing the forest for the trees: hierarchical displays of hypertext structures. *ACM Office Information Systems '88*, 205-212.

[Fisc+91] Fischer, G. & C. Stevens (1991). Information access in complex, poorly structured information spaces. *CHI '91, Human Factors in Computing Systems*, 63-70.

[Fitt54] Fitts, P.M. (1954). The information capacity of the human motor system in controlling amplitude of movement. *Journal of Experimental Psychology, 47*, 381-391.

[Flem98] Fleming, J. (1998). *Web Navigation: Designing the User Experience*. Sebastopol, CA: O'Reilly.

[Freu01] Freudenthal, D. (2001). Age differences in the performance of information retrieval tasks. *Behaviour and Information Technology, 20(1)*, 9-22.

[Frøk+00] Frøkjær, E., M. Hertzum, & K. Hornbæk (2000). Measuring usability: are effectiveness, efficiency, and satisfaction really correlated? *CHI '00, Human Factors in Computing Systems*, 345-352.

[Furn81] Furnas, G.W. (1981). The FISHEYE view: a new look at structured files. Technical Report 81-11221-9, Bell Laboratories.

[Furn86] Furnas, G.W. (1986). Generalized fisheye views. *CHI '86, Human Factors in Computing Systems*, 16-23.

[Furn97] Furnas, G.W. (1997). Effective view navigation. *CHI '97, Human Factors in Computing Systems*, 367-374.

[Furn+87] Furnas, G.W., T.K. Landauer, L.M. Gomez, & S.T. Dumais (1987). The vocabulary problem in human-system communication. *Communications of the ACM, 30(11)*, 964-971.

[Furu+90] Furuta, R. & P.D. Stotts (1990). A functional meta-structure for hypertext models and systems. *Electronic Publishing, 3(4)*, 179-205.

[Furu+97] Furuta, R., F.M. Shipman III, C.C. Marshall, D. Brenner, & H. Hsieh (1997). Hypertext paths and the World-Wide Web: experiences with Walden's Paths. *ACM Hypertext '97*, 167-176.

[Gee01] Gee, K. (2001). The ergonomics of hypertext narrative: usability testing as a tool for evaluation and redesign. *ACM Journal of Computer Documentation, 25(1)*, 3-16.

[Germ+00] Germán, D.M. & D.D. Cowan (2000). Towards a unified catalog of hypermedia design patterns. *IEEE Proceedings of the 33rd Hawaiian International Conference on System Sciences*.

[Gill+00] Gillenson, M., D.L. Sherrell, & L. Chen (2000). A taxonomy of Web site traversal patterns and structures. *Communications of the Association for Information Systems, 3*, Article No. 17.

[Gree+99] Greenberg, S., & A. Cockburn (1999). Getting Back to Back: alternative behaviors for a Web browser's Back button. *Proceedings of the Fifth Annual Human Factors and the Web Conference*.

[Gray90] Gray, S.H. (1990). Using protocol analyses and drawings to study mental model construction during hypertext navigation. *International Journal of Human-Computer Interaction, 2*, 359-377.

[Grud89] Grudin, J. (1989). The case against user interface consistency. *Communications of the ACM, 32(10)*, 1164-1173.

[Haas+98] Haas, S.W. & E.S. Grams (1998). Page and link classifications: connecting diverse resources. *ACM Digital Libraries '98*, 99-107.

[Hala87] Halasz, F.G. (1987). Reflections on Notecards: seven issues for the next generation of hypermedia systems. *ACM Hypertext '87*, 345-365.

[Hala+83] Halasz, F.G. & T.P. Moran (1983). Mental models and problem solving in using a calculator. *CHI '83, Human Factors in Computing Systems*, 212-216.

[Hara59] Harary, F. (1959). Status and contrastatus. *Sociometry, 22*, 23-43.

[Hard89] Hardman, L. (1989). Evaluating the usability of the Glasgow Online hypertext. *Hypermedia, 1*, 34-63.

[Hend89] Hendrickson, J.J. (1989). Performance preference and visual scan patterns on a menu-based system: implications for interaction design. *CHI '89, Human Factors in Computing Systems*, 217-222.

[Hirt+98] Hirtle, S.C., M.E. Sorrows, & G. Cai (1998). Clusters on the World Wide Web: creating neighborhoods of make-believe. *ACM Hypertext '98*, 289-290.

[Hoch+00] Hochheiser, H. & B. Shneiderman (2000). Performance benefits of simultaneous over sequential menus as task complexity increases. *International Journal of Human-Computer Interaction, 12(2)*, 173-192.

[Hofm+99] Hofman, R. & H. van Oostendorp (1999). Cognitive effects of a structural overview in a hypertext. *British Journal of Educational Technology, 30(2)*, 129-140.

[Höök+96] Höök, K., M. Sjölinder & N. Dahlbäck (1996). Individual differences and navigation in hypermedia. *Eighth European Conference on Cognitive Ergonomics*.

[Horn+97] Hornof, A.J. & D.E. Kieras (1997). Cognitive modeling reveals menu search is both random and systematic. *CHI '97, Human Factors in Computing Systems*, 107-114.

[Horn+99] Hornof, A.J. & D.E. Kieras (1999). Cognitive modeling demonstrates how people use anticipated location knowledge of menu Items. *CHI '99, Human Factors in Computing Systems*, 410-417.

[Howe94] Howes, A. (1994). A model of the acquisition of menu knowledge by exploration. *CHI '94, Human Factors in Computing Systems*, 445-451.

[Howe+91] Howes, A. & R.M. Young (1991). Predicting the learnability of task-action mappings. *CHI '91, Human Factors in Computing Systems*, 113-118.

[Hube+98] Huberman, B.A., P. Pirolli, J. Pitkow, & R.M. Lukose (1998). Strong regularities in World Wide Web surfing. *Science, 280(5360)*, 95-97.

[Ido+96] Ido, K. & T. Yamaoka (1996). Examining basic items of a screen design. *CHI '96, Human Factors in Computing Systems*, 79-80.

[Inst+93] Instone, K., B.N. Teasley, & L.M. Leventhal (1993). Empirically-based re-design of a hypertext encyclopedia. *CHI '93, Human Factors in Computing Systems*, 500-506.

[Ivor+01] Ivory, M.Y., R.R. Sinha, & M.A. Hearst (2001). Empirically validated Web page design metrics. *CHI '01, Human Factors in Computing Systems*, 53-60.

[John83] Johnson-Laird, P.N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge, MA: Harvard.

[John88] Johnson-Laird, P.N. (1988). *The Computer and the Mind: An Introduction to Cognitive Science*. Cambridge, MA: Harvard.

[John91] Johnson-Laird, P.N. (1991). Mental Models. In M.I. Posner (Ed.), *Foundations of Cognitive Science*, Cambridge, MA: MIT, 469-499.

[Jone+96] Jones, S. & A. Cockburn (1996). A study of navigational support provided by two World Wide Web browsing applications. *ACM Hypertext '96*, 161-169.

[Jul+97] Jul, S. & G.W. Furnas (1997). Navigation in electronic worlds. *SIGCHI Bulletin, 29(4)*, 44-49.

[Kaas+01] Kaasten, S., S. Greenberg, & C. Edwards (2001). How people recognize previously seen Web pages from titles, URLs, and thumbnails. Technical Report 2001-692-15, Department of Computer Science, University of Calgary.

[Kahn95] Kahn, P. (1995). Visual cues for local and global coherence in the WWW. *Communications of the ACM, 38(8)*, 67-69.

[Kapt93] Kaptelinin, V. (1993). Item recognition in menu selection: the effect of practice. *INTERACT '93 and CHI '93 Conference Companion on Human Factors in Computing Systems*, 183-184.

[Kemp+99] Kemp, B. & K. Buckner (1999). A taxonomy of design guidance for hypermedia design. *Interacting with Computers, 12(2)*, 143-160.

[Keye+92] Keyes, E. & R. Krull (1992). User information processing strategies and online visual structure. *ACM Systems Documentation '92*, 121-128.

[KimH+95] Kim, H. & S. Hirtle (1995). Spatial metaphors and disorientation in hypertext browsing. *Behavior and Information Technology, 14*, 239-250.

[KimJ99] Kim, J. (1999). An empirical study of navigation aids in customer interfaces. *Behaviour and Information Technology, 18(3)*, 213-224.

[KimJ+98] Kim, J. & B. Yoo (1998). Does every link have the same usability? An exploratory study of the link structure of cyber malls. *CHI '98, Human Factors in Computing Systems*, 315-316.

[Koen+96] Koenemann, J. & N.J. Belkin (1996). A case for interaction: a study of interactive information retrieval behaviour and effectiveness. *CHI '96, Human Factors in Computing Systems*, 205-212.

[Lars+98] Larson, K. & M. Czerwinski (1998). Web page design: implications of memory, structure, and scent for information retrieval. *CHI '98, Human Factors in Computing Systems*, 25-32.

[Lesk97] Lesk, M. (1997). *Practical Digital Libraries: Books, Bytes, and Bucks*. San Francisco, CA: Morgan Kaufmann.

[Lewi+99] Lewis, P.H., W. Hall, L.A. Carr, & D. De Roure (1999). The significance of linking. *ACM Computing Surveys, 31(4es)*, Article No. 10.

[Lin+00] Lin, J., M.W. Newman, J.I. Hong, & J.A. Landay (2000). DENIM: finding a tighter fit between tools and practice for Web site design. *CHI '00, Human Factors in Computing Systems*, 510-517.

[Liu+94] Liu, M. & M.W. Reed (1994). The relationship between the learning strategies and learning styles in a hypermedia environment. *Computers in Human Behavior, 10(4)*, 419-434.

[Luko+98] Lukose, R.M. & B.A. Huberman (1998). Surfing as a real option. *ACM Information and Computation Economies '98*, 45-51.

[LyncK60] Lynch, K. (1960). *The Image of the City*. Cambridge, MA: M.I.T. Press.

[LyncP+97] Lynch, P.J. & S. Horton (1997). Imprudent linking weaves a tangled Web. *IEEE Computer, 30*, 115-117.

[Mack+91] Mackinlay, J.D., G.G. Robertson, & S.K. Card (1991). The Perspective Wall: detail and context smoothly integrated. *CHI '91, Human Factors in Computing Systems*, 173-179.

[Macs+98] Macskassy, S.A., A. Banerjee, B.D. Davison, & H. Hirsh (1998). Human performance on clustering Web pages. *Proceedings of the Fourth International Conference on Knowledge Discovery and Data Mining*, 264-268.

[Magl+97] Maglio, P.P. & R. Barrett (1997). On the trail of information searchers. *The 19th Annual Meeting of the Cognitive Science Society*, University of California, San Diego.

[Magl+98] Maglio, P.P. & T. Matlock (1998). Metaphors we surf the Web by. Workshop on Personalized and Social Navigation in Information Space, Stockholm, Sweden.

[Magl+00] Maglio, P.P. & C.S. Campbell (2000). Tradeoffs in displaying peripheral information. *CHI '00, Human Factors in Computing Systems*, 241-248.

[Marc95] Marchionini, G.M. (1995). *Information Seeking in Electronic Environments*. Cambridge, England: Cambridge.

[Marc+88] Marchionini, G. & B. Shneiderman (1988). Finding facts vs. browsing knowledge in hypertext systems. *IEEE Computer, 21(1)*, 70-80.

[Mars+89] Marshall, C.C. & P.M. Irish (1989). Guided tours and on-line presentations: how authors make hypertext intelligible for readers. *ACM Hypertext '89*, 15-26.

[Mars+93] Marshall, C.C. & F.M. Shipman III (1993). Searching for the missing link: discovering implicit structure in spatial hypertext. *ACM Hypertext '93*, 217-230.

[Mars+95] Marshall, C.C. & F.M. Shipman III (1995). Spatial hypertext: designing for change. *Communications of the ACM, 38(8)*, 88-97.

[Matl+96] Matlock, T. & P.P. Maglio (1996). Apparent motion on the World Wide Web. *The 18th Annual Meeting of the Cognitive Science Society*, University of California, San Diego.

[McDo+96] McDonald, S. & R.J. Stevenson (1996). Disorientation in hypertext: the effects of three text structures on navigation performance. *Applied Ergonomics, 27(1)*, 61-68.

[McDo+98] McDonald, S. & R.J. Stevenson (1998). 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(2)*, 129-142.

[McEn99] McEneaney, J.E. (1999). Visualizing and assessing navigation in hypertext. *ACM Hypertext '99*, 61-70.

[McEn00] McEneaney, J.E. (2000). Navigational correlates of comprehension in hypertext. *ACM Hypertext '00*, 254-255.

[McKe+01] McKenzie, B. & A. Cockburn (2001). An empirical analysis of Web page revisitation. *IEEE Proceedings of the 34th Hawaiian International Conference on System Sciences*.

[McKn+91] McKnight, C., A. Dillon, & J. Richardson (1991). *Hypertext in Context*. Cambridge, England: Cambridge.

[Mead+97] Mead, S.E., V.A. Spaulding, R.A. Sit, B. Meyer, & N. Walker (1997). Effects of age and training on World Wide Web navigation strategies. *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*.

[Mehl92] Mehlenbacher, B. (1992). Navigating online information: a characterization of extralinguistic factors that influence user behavior. *ACM Systems Documentation '92*, 35-46.

[Meye+97] Meyer, B., R.A. Sit, V.A. Spaulding, S.E. Mead, & N. Walker (1997). Age group differences in World Wide Web navigation. *CHI '97, Human Factors in Computing Systems, 2*, 295-296.

[Mill56] Miller, G.A. (1956). The magic number seven, plus or minus two: some limits on our capacity for processing information. *The Psychological Review, 63*, 81-97.

[Mitc+95] Mitchell, R., D. Day, & L. Hirschman (1995). Fishing for information on the internet. *IEEE Information Visualization '95*, 105-111.

[Mizz97] Mizzaro, S. (1997). Relevance: the whole history. *Journal of the American Society for Information Science, 48(9)*, 810-832.

[Mizz98] Mizzaro, S. (1998). How many relevances in information retrieval? *Interacting with Computers, 10(3)*, 303-320.

[Modj+97] Modjeska, D. & A. Marsh (1997). Structure and memorability of Web sites. Technical Report, Computer Systems Research Group, University of Toronto.

[Moha92] Mohageg, M.F. (1992). The influence of hypertext linking structures on the efficiency of information retrieval. *Human Factors, 34(3)*, 351-367.

[Mora98] Moran, T.P. (1998). Designing usable lists. *ACM Systems Documentation '98*, 58-62.

[More+98] Moreau, L. & W. Hall (1998). On the expressiveness of links in hypertext systems. *The Computer Journal, 41(7)*, 459-473.

[Morr+00] Morrison, J.B., P. Pirolli, & S.K. Card (2000). A taxonomic analysis of what World Wide Web activities significantly impact people's decisions and actions. Technical report UIR-R-2000-17, Xerox PARC.

[Mukh+95a] Mukherjea, S. & J.D. Foley (1995). Showing the context of nodes in the World Wide Web. *CHI '95, Human Factors in Computing Systems*, 326-327.

[Mukh+95b] Mukherjea, S., J.D. Foley, & S. Hudson (1995). Visualizing complex hypermedia networks through multiple hierarchical views. *CHI '95, Human Factors in Computing Systems*, 331-337.

[Mukh+97] Mukherjea, S. & H. Yoshinori (1997). Focus+Context views of World-Wide Web nodes. *ACM Hypertext '97*, 187-196.

[Meye+97] Meyer, B.A., R.A. Sit, V.A. Spaulding, S.E. Mead, & N. Walker (1997). Age group differences in World Wide Web navigation. *CHI '97, Human Factors in Computing Systems*, 295-296.

[Nana+99] Nanard, J. & M. Nanard (1999). Toward an hypermedia design patterns space. Second Workshop in Hypermedia Development: Design Patterns in Hypermedia.

[Nati+97] Nation, D.A., C. Plaisant, G. Marchionini, & A. Komlodi (1997). Visualizing Websites using a hierarchical table of contents browser: WebTOC. *Proceedings of the Third Annual Human Factors and the Web Conference*.

[Nava+99] Navarro-Prieto, R., M. Scaife, & Y. Rogers (1999). Cognitive strategies in Web searching. *Proceedings of the Fifth Annual Human Factors and the Web Conference*.

[Neer+01] Neerincx, M.A., J. Lindenberg, & S. Pemberton (2001). Support concepts for Web navigation: a cognitive engineering approach. *ACM International World Wide Web Conference '01*, 119-128.

[Newe90] Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard.

[Newf+98] Newfield, D., B.S. Sethi, & K. Ryall (1998). Scratchpad: mechanisms for better navigation in directed Web searching. *ACM User Interface Software and Technology '98*, 1-8.

[Newm+00] Newman, M.W. & J.A. Landay (2000). Sitemaps, storyboards, and specifications: a sketch of Web site design practice. *ACM Designing Interactive Systems '00*, 263-274.

[Ngo01] Ngo, D.C.L. (2001). Measuring the aesthetic elements of screen designs. *Displays, 22*, 73-78.

[Ngo+00] Ngo, D.C.L., L.S. Teo, & J.G. Byrne (2000). Formalising guidelines for the design of screen layouts. *Displays, 21*, 3-15.

[Niel89] Nielsen, J. (1989). The matters that really matter for hypertext usability. *ACM Hypertext '89*, 239-248.

[Niel90a] Nielsen, J. (1990). The art of navigating through hypertext. *Communications of the ACM, 33(3)*, 296-310.

[Niel90b] Nielsen, J. (1990). *Hypertext and Hypermedia*. San Diego, CA: Academic.

[Niel93] Nielsen, J. (1993). *Usability Engineering*. San Francisco, CA: Morgan Kaufmann.

[Niel94] Nielsen, J. (1994). Enhancing the explanatory power of usability heuristics. *ACM CHI '94, Human Factors in Computing Systems*, 152-158.

[Niel95] Nielsen, J. (1995). *Multimedia and Hypertext: The Internet and Beyond*. Boston, MA: AP Professional.

[Niel99] Nielsen, J. (1999). User interface directions for the Web. *Communications of the ACM, 42(1)*, 65-72.

[Niel00] Nielsen, J. (2000). *Designing Web Usability: The Practice of Simplicity*. Indianapolis, IN: New Riders.

[Niel+94] Nielsen, J. & J. Levy (1994). Measuring usability preference vs. performance. *Communications of the ACM, 37(4)*, 66-75.

[Niel+96] Nielsen, J. & A. Wagner (1996). User interface design for the WWW. *CHI '96 Conference Companion on Human Factors in Computing Systems*, 330-331.

[Niev+80] Nievergelt, J. & J. Weydert (1980). Sites, modes, and trails: telling the user of an interactive system how to get to places. In R.A. Guedj, P.J.W. ten Hagen, F.R.A. Hopgood, H.A. Tucker, & D.A. Duce (Eds.), *Methodology of Interaction*. Amsterdam: North Holland, 327-338.

[Niga+98] Nigay, L. & F. Vernier (1998). Design method of interaction techniques for large information spaces. *ACM Advanced Visual Interfaces '98*, 37-46.

[NormD83] Norman, D.A. (1983). Design principles for human-computer interfaces. *CHI '83, Human Factors in Computing Systems*, 1-10.

[NormK91] Norman, K.L. (1991). *The Psychology of Menu Selection: Designing Cognitive Control at the Human/Computer Interface*. Norwood, NJ: Ablex.

[Nort+97] North, C. & B. Shneiderman (1997). A taxonomy of multiple window coordinations. Technical Report CS-TR-3854, University of Maryland Computer Science Department.

[O'Day+93] O'Day, V.L. & R. Jeffries (1993). Orienteering in an information landscape: how information seekers get from here to there. *CHI '93, Human Factors in Computing Systems*, 438-445.

[Øest99] Øesterbye, K. (1999). Hierarchical structure through navigation side bars. Second Workshop in Hypermedia Development: Design Patterns in Hypermedia.

[Otte+00] Otter, M. & H. Johnson (2000). Lost in hyperspace: metrics and mental models. *Interacting with Computers, 13(1)*, 1-40.

[Park+00a] Park, J. & J. Kim (2000). Contextual navigation aids for two World Wide Web systems. *International Journal of Human-Computer Interaction, 12(2)*, 193-217.

[Park+00b] Park, J. & J. Kim (2000). Effects of contextual navigation aids on browsing diverse Web systems. *CHI '00, Human Factors in Computing Systems*, 257-264.

[Paru89] Parunak, H.V.D. (1989). Hypermedia topologies and user navigation. *ACM Hypertext '89*, 43-50.

[Pilg+99] Pilgrim, C.J. & Y.K. Leung (1999). Designing WWW site map systems. *IEEE Proceedings of the Tenth International Workshop on Database and Expert Systems*, 253-258.

[Piro+95] Pirolli, P. & S.K. Card (1995). Information foraging in information access environments. *CHI '95, Human Factors in Computing Systems*, 51-58.

[Piro+96] Pirolli, P., J. Pitkow, & R. Rao (1996). Silk from a sow's ear: extracting usable structures from the Web. *CHI '96, Human Factors in Computing Systems*, 118-125.

[Piro+99] Pirolli, P. & S.K. Card (1999). Information Foraging. *Psychological Review, 106(4)*, 643-675.

[Piro+00] Pirolli, P., S.K. Card, & M.M. Van Der Wege (2000). The effect of information scent on searching information visualizations of large tree structures. *ACM Advanced Visual Interfaces '00*, 161-172.

[Pitk98] Pitkow, J. (1998). Summary of WWW characterizations. *Seventh International World Wide Web Conference*.

[Pitk+96] Pitkow, J. & C.M. Kehoe (1996). Emerging trends in the WWW user population. *Communications of the ACM, 39(6)*, 106-108.

[Pitk+97] Pitkow, J. & P. Pirolli (1997). Life, death, and lawfulness on the electronic frontier. *CHI '97, Human Factors in Computing Systems*, 383-390.

[Pitk+99] Pitkow, J. & P. Pirolli (1999). Mining longest repeated subsequences to predict World Wide Web surfing. *Second USENIX Symposium on Internet Technologies and System*.

[Rada+92] Rada, R. & C. Murphy (1992). Searching versus browsing in hypertext. *Hypermedia, 4(1)*, 1-30.

[Rao+94] Rao, R. & S.K. Card (1994). The Table Lens: merging graphical and symbolic representations in an interactive Focus+Context visualization for tabular information. *CHI '94, Human Factors in Computing Systems*, 318-322.

[Ratn+96] Ratner, J., E.M. Grose, & C. Forsythe (1996). Characterization and assessment of HTML style guides. *CHI '96, Human Factors in Computing Systems*, 115-116.

[Reic+99] Reich, S., L. Carr, D. De Roure, & W. Hall (1999). Where have you been from here? Trials in hypertext systems. *ACM Computing Surveys, 31(4es)*, Article No. 11.

[Reis99] Reisman, J. (1999). Web site design: less is more. *IEEE IT Professional, 1(5)*, 63-64.

[Rivl+94] Rivlin, E., R. Botafogo, & B. Shneiderman (1994). Navigating in hyperspace: designing a structure-based toolbox. *Communications of the ACM, 37(2)*, 87-96.

[Robe+81] Robertson, G.G., D. McCracken, & A. Newell (1981). The ZOG approach to man-machine communications. *International Journal of Man-Machine Studies, 14*, 461-488.

[Roue92] Rouet, J.F. (1992). Cognitive processing of hyperdocuments: when does nonlinearity help? *ACM Hypertext '92*, 131-140.

[Ruge95] Ruge, G. (1995). Human memory models and term association. *ACM Research and Development in Information Retrieval '95*, 219-227.

[Rump98] Rumpradit, C. (1998). From novices to experts: user performance, confidence, and satisfaction on the World Wide Web. *Association for Information Systems, Fourth Americas Conference*.

[Rump+99] Rumpradit, C. & M.L. Donnell (1999). Navigational cues on user interface design to produce better information seeking on the World Wide Web. *IEEE Proceedings of the 32nd Hawaiian International Conference on System Sciences*.

[Scha+96] Schaffer, D., Z. Zuo, S. Greenberg, L. Bartram, J. Dill, S. Dubs, & M. Roseman (1996). Navigating hierarchically clustered networks through fisheye and full-zoom methods. *ACM Transactions on Computer-Human Interaction, 3(2)*, 162-188.

[Sear93] Sears, A. (1993). Layout appropriateness: a metric for evaluating user interface widget layout. *IEEE Transactions on Software Engineering, 19(7)*, 707-719.

[Sear+00] Sears, A. & J.A. Jacko (2000). Understanding the relation between network quality of service and the usability of distributed multimedia documents. *Human-Computer Interaction, 15(1)*, 43-68.

[Shne87] Shneiderman, B. (1987). User interface design and evaluation for an electronic encyclopedia. In G. Salvendy, S.L. Sauter, & J.J. Hurell Jr. (Eds.), *Social, Ergonomic, and Stress Aspects of Work with Computers*. Amsterdam: Elsevier Science, 113-122.

[Shne96] Shneiderman, B. (1996). The eyes have it: a task by data type taxonomy for information visualizations. *IEEE Visual Languages '96*, 336-343.

[Shne97] Shneiderman, B. (1997). Designing information-abundant Web sites: issues and recommendations. *International Journal of Human-Computer Studies, 47*, 5-29.

[Shne98] Shneiderman, B. (1998). *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Reading, MA: Addison-Wesley.

[Shne+98] Shneiderman, B., J. Nielsen, S. Butler, M. Levi, & F. Conrad (1998). Is the Web really different from everything else? *CHI '98 Summary: Human Factors in Computing Systems*, 92-93.

[Shub+97] Shubin, H. & M.M. Meehan (1997). Navigation in Web applications. *ACM Interactions, 4(6)*, 13-17.

[Shum90] Shum, S.B. (1990). Real and virtual spaces: mapping from spatial cognition to hypertext. *Hypermedia, 2(2)*, 133-158.

[Simp+89] Simpson, A. & C. McKnight (1989). Navigation in hypertext: structural cues and mental maps. In R. McAleese & C. Green (Eds.), *Hypertext: State of the Art*. Oxford, England: Intellect Books, 73-83.

[Smit96] Smith, P.A. (1996). Towards a practical measure of hypertext usability. *Interacting with Computers, 8(4)*, 365-381.

[Smit+97] Smith, P.A., I.A. Newman, & L.M. Parks (1997). Virtual hierarchies and virtual networks: some lessons from hypermedia usability research applied to the World Wide Web. *International Journal of Human-Computer Studies, 47*, 67-95.

[Stan+92] Stanton, N., R. Taylor, & L. Tweedie (1992). Maps as navigational aids in hypertext environments: an empirical evaluation. *Journal of Educational Multimedia and Hypermedia, 1(4)*, 431-444.

[Stev+78] Stevens, A. & P. Coupe (1978). Distortions in judged spatial relations. *Cognitive Psychology, 10*, 422-437.

[Stre+85] Streeter, L., D. Vitello, & S. Wonsiewicz (1985). How to tell people where to go: comparing navigational aids. *International Journal of Man-Machine Studies, 22*, 549-562.

[Sutc+98] Sutcliffe, A. & M. Ennis (1998). Towards a cognitive theory of information retrieval. *Interacting with Computers, 10(3)*, 321-351.

[Taba+98] Tabatabai, D. & L. Francesca (1998). Expert-novice differences in searching the Web. *Association for Information Systems, Fourth Americas Conference*.

[Taus+97a] Tauscher, L.M. & S. Greenberg (1997). How people revisit Web pages: empirical findings and implications for the design of history mechanisms. *International Journal of Human-Computer Studies, 47*, 94-137.

[Taus+97b] Tauscher, L.M. & S. Greenberg (1997). Revisitation patterns in World Wide Web navigation. *CHI '97, Human Factors in Computing Systems*, 399-406.

[Teit+83] Teitelbaum, R.C. & R.E. Granda (1983). The effects of positional constancy on searching menus for information. *CHI '83, Human Factors in Computing Systems*, 150-153.

[Then+95] Theng, Y.L., H. Thimbleby & M. Jones (1995). Reducing information overload: a comparative study of hypertext systems. IEE Colloquium: Information Overload, 6/1-6/4.

[Thür+91] Thüring, M., J.M. Haake, & J. Hannemann (1991). What's Eliza doing in the Chinese Room? Incoherent hyperdocuments – and how to avoid them. *ACM Hypertext '91*, 161-177.

[Thür+95] Thüring, M., J. Hannemann, & J.M. Haake (1995). Hypermedia and cognition: designing for comprehension. *Communications of the ACM, 38(8)*, 57-66.

[Thur98] Thury, E.M. (1998). Analysis of student Web browsing behavior: implications for designing and evaluating Web sites. *ACM Systems Documentation '98*, 265-270.

[Tolm48] Tolman, E.C. (1948). Cognitive maps in rats and men. *The Psychological Review, 55(4)*, 189-208.

[Tosc00] Tosca, S.P. (2000). A pragmatics of links. *ACM Hypertext '00*, 77-84.

[Trac+00] Tractinsky, N., A.S. Katz, & D. Ikar (2000). What is beautiful is usable. *Interacting with Computers, 13(2)*, 127-145.

[Trav89] Travers, M. (1989). A visual representation for knowledge structures. *ACM Hypertext '89*, 147-158.

[Trie+80] Triesman, A. & G. Gelade (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97-136.

[Trig+86] Trigg, R.H. & M. Weiser (1986). TEXTNET: a network-based approach to text handling. *ACM Transactions on Information Systems, 4(1)*, 1-23.

[Turi+97] Turine, M.A.S., M. Oliveira, & P.C.A. Masiero (1997). A navigation-oriented hypertext model based on statecharts. *ACM Hypertext '97*, 102-111.

[Turo+92] Turo, D. & B. Johnson (1992). Improving the visualization of hierarchies with treemaps: design issues and experimentation. *IEEE Visualization '92*, 124-131.

[Utti+89] Utting, K. & N. Yankelovich (1989). Context and orientation in hypermedia networks. *ACM Transactions on Information Systems, 7(1)*, 58-84.

[Vald+88] Valdez, F. & M. Chignell (1988). Browsing models for hypermedia databases. *Proceedings of the Human Factors Society 32nd Annual Meeting*, 318-322.

[vanN+99] van Nimwegen, C., M. Pouw, & H. van Oostendorp (1999). The influence of structure and reading-manipulation on usability of hypertexts. *Interacting with Computers, 12(1)*, 7-21.

[vanS+01] van Schaik, P. & J. Ling (2001). The effects of frame layout and differential background contrast on visual search performance in Web pages. *Interacting with Computers, 13(5)*, 513-525.

[vanW98] van Waes, L. (1998). Evaluating on-line and off-line searching behavior using thinking-aloud protocols to detect navigation barriers. *ACM Systems Documentation '98*, 180-183.

[Vann99] Vannucchi de Almeida Santos, H. (1999). Thwarting the Web users' expectations. *ACM Creativity and Cognition '99*, 199-200.

[Vice+88] Vicente, K.J. & R.C. Williges (1988). Accommodating individual differences in searching a hierarchical file system. *International Journal of Man-Machine Studies, 29*, 647-668.

[Vins99] Vinson, N.G. (1999). Design guidelines for landmarks to support navigation in virtual environments. *CHI '99, Human Factors in Computing Systems*, 278-285.

[Vora+94] Vora, P.R., M.G. Helander, & V.L. Shalin (1994). Evaluating the influence of interface styles and multiple access paths in hypertext. *CHI '94, Human Factors in Computing Systems*, 323-329.

[Whis+95] Whisenand, T.G. & H.H. Emurian (1995). Some effects of angle of approach on icon selection. *CHI '95, Human Factors in Computing Systems*, 298-299.

[Winc+01] Winckler, M., C. Farenc, P. Palanque, & R. Bastide (2001). Designing navigation for Web interfaces. *IHM Conference on Human-Computer Interaction '01*.

[Wood84] Woods, D.D. (1984). Visual momentum: a concept to improve the cognitive coupling of person and computer. *International Journal of Man-Machine Studies, 21*, 229-244.

[Wrig91] Wright, P. (1991). Cognitive overheads and prostheses: some issues in evaluating hypertexts. *ACM Hypertext '91*, 1-12.

[Wrig+90] Wright, P. & A. Lickorish (1990). An empirical comparison of two navigation systems for two hypertexts. In R. McAleese & C. Green (Eds.), *Hypertext: State of the Art*. Oxford, England: Intellect Books, 84-93.

[Yama+95] Yamada, S., J. Hong, & S. Sugita (1995). Development and evaluation of hypermedia for museum education: validation of metrics. *ACM Transactions on Computer-Human Interaction, 2(4)*, 284-307.

[Yu+01] Yu, B., S. Roh, & S. Han (2001). The effect of different menu styles on the user's perception and performance on the WWW. *Proceedings of Topics on Distance Learning '01*, Purdue University.

[Zaph+00] Zaphiris, P. & D. Ellis (2000). Mathematical modeling of age differences in hierarchical navigation systems. *ACM Universal Usability '00*, 157-158.

[Zell+98] Zellweger, P.T., B. Chang, & J.D. Mackinlay (1998). Fluid links for informed and incremental link transitions. *ACM Hypertext '98*, 50-57.

[ZhanH+01] Zhang, H. & G. Salvendy (2001). The implications of visualization ability and structure preview design for Web information search tasks. *International Journal of Human-Computer Interaction, 13(1)*, 75-95.

[ZhanX+00] Zhang, X., K.B. Keeling, & R.J. Pavur (2000). Information quality of commercial Web site home pages: an explorative analysis. *ACM Hypertext '98*, 164-175.

[Zimm+00] Zimmerman, D. & P. Walls (2000). Exploring navigational patterns on the Web. *IEEE 18th Annual Conference on Computer Documentation*, 581-591.

[Zipf49] Zipf, G. (1949). *Human Behavior and the Principle of Least Effort*. Cambridge, MA: Addison-Wesley.

Acknowledgments

Study in Human-Computer Interaction is, by its very nature, interdisciplinary. Without a solid grounding in fundamental aspects of both the human and the computer, one cannot possibly hope to truly understand the interaction. This personal belief has made Stanford University, and the Symbolic Systems Program in particular, the perfect intellectual environment for my study, research, and growth over the recent years. I have benefited greatly from the diversity of knowledge available to me at Stanford, and know that I have barely scratched the surface of it.

Numerous individuals have contributed suggestions, thoughts, references, potential problems, and perspectives that have shaped the process resulting in this thesis, with a few having significant impact in the long run. As this process began with my undergraduate work, I must acknowledge that the Symbolic Systems Honors Program, and its associated courses (including the Honors College), were especially critical early in the process. Todd Davies and fellow students made useful comments throughout that critical period, and taught me valuable lessons about the development and investigation of ideas.

I am greatly appreciative of the assistance and feedback of Clifford Nass, whose extensive experimental expertise has benefited my work substantially, and has made me better appreciate that no experimental design issue can be taken for granted or ignored.

Finally, I must acknowledge my adviser, Terry Winograd, for his candidness and willingness to advise my work in the first place. The impact of my conversations with him over the past two years is significant throughout this thesis. There undoubtedly remain ideas I have neglected, issues I have treated glibly, and perspectives I have not considered, but their incidence is greatly reduced due to his valuable feedback throughout the process. More important, it is obvious to me that the very way I approach and think about this field and its major issues has been heavily influenced by that interaction.