Metaphors We Surf the Web By[1]

Paul P. Maglio
IBM Almaden Research Center
650 Harry Rd., NWED-B2
San Jose, CA 95120
pmaglio@almaden.ibm.com

Teenie Matlock
Cognitive Psychology Program
University of California, Santa Cruz
Santa Cruz, CA 95064
tmatlock@cats.ucsc.edu

Abstract

The way people think about the World-Wide Web (WWW) has implications for the way that they navigate it. In this paper, we discuss the nature of people's *metaphorical* conception of the WWW, as gathered from interviews with beginning and experienced web users. Based on linguistic data, we argue that people naturally think of the web as a kind of physical space in which they move, although information on the web is not physical, and web users do not actually move. Nevertheless, such metaphorical thought is motivated by the same basic image schemata that people rely on to mentally structure everyday life.

Introduction

How do people *naturally* conceive of information spaces, such as the World-Wide Web? It seems obvious that the way people think about the web would have implications for the way they use and navigate the web, and thus for the design of tools to facilitate information navigation. In this paper, we consider metaphorical conceptions of the WWW, as gathered primarily from language-use data. Briefly, we argue that both experienced and inexperienced web users naturally talk about the web in definite and consistent ways. For instance, people see themselves as moving toward information, rather than information as moving toward them. Nevertheless, we found some differences between experienced and inexperienced web users in the consistency with which they talked about web activities. In the end, we argue that the particular language people use is metaphorical and is motivated by basic image schemata, which emerge from embodied experience (e.g., Johnson, 1987; Lakoff, 1987).

Before presenting our data and argument in detail, we first discuss some prior research concerning the way people think about the web, along with some background on metaphor and thought.

How People Remember the Web

Based on data collected from people asked to recall specific WWW searches, Maglio and Barrett (1997a) argued that web navigation is conceived in terms of a cognitive map similar to a cognitive map of physical space, that is, in terms of landmarks and routes (e.g., Anderson, 1980). In this study, experienced users searched the web for answers to specific questions. To identify key cognitive aspects of their activities, users were first asked about their plans, and then their behavior was tracked while they searched. Then a day later they were asked to recall the steps they had taken in each of their searches the previous day, and finally to retrace their steps. Participants were not warned on the first day that recall would be required on the second day. This method enabled Maglio and Barrett both to chart behavior to uncover search tactics (using the behavioral traces) and to extract some of the structure of their internal representations (using the recall data).

The data showed that participants recalled only a few of the sites they visited. Specifically, they remembered key nodes that led to the target information. These nodes were called *anchor points* by analogy to the notion of anchor points in the cognitive map literature (Couclelis, Golledge, Gale & Tobler, 1987). An anchor was defined as a node along a search path from which there is an unbroken sequence of links on successive pages that lead to the goal node (i.e., no URLs need to be typed in or explicitly recalled). Once traversed, anchor points are recognized as lying along the path to the goal—even if the same path is not followed to the goal in every case. For the participants in the study, searching on the second

[1] With apologies to Lakoff and Johnson (1980).

day often meant finding anchors encountered on the first day, rather than finding paths found on the first day.

A second observation that emerged from the behavioral data is that individuals relied on personal routines when trying to find information. For instance, some participants routinely used a particular search engine, such as AltaVista, whereas others routinely used a particular hierarchical catalog, such as Yahoo! It is not merely that these searchers preferred to use one approach over another, but that they conceptualized their search tasks in terms of their favorite routines. It often did not matter what was actually done on the first day, the searchers remembered searching as if their personal routines had been followed. On the analogy to cognitive maps of physical space, personal routines correspond to the familiar routes that an individual uses to get from one landmark (or anchor point) to another.

If people mentally structure web use in this way, tools for web navigation ought to present the web in this way. Because individuals tend to use the same search patterns over and over, and because they recall their searches in terms of their standard patterns—almost regardless of what they actually did—Maglio and Barrett (1997b) built a personal web agent to identify repeated search patterns and to suggest similar patterns for new searches. Because people focus on key nodes or anchor points when recalling their searches, and because these structure memory for the searches, Maglio and Barrett (1997b) built a web agent to identify the key nodes in finding a piece of information, and to maintain personal trails in terms of these.

How People Talk About the Web

The key to designing information navigation tools lies in discovering how people naturally conceive of information spaces. Technically, the WWW is part of a network of geographically distributed machines connected via wires. The information accessible by users of this physical network is organized in a conceptual network of hyperlinks among documents. Despite this actual structure, people's conceptual structure of the WWW is rather different.

Matlock and Maglio (1996) found that web users often refer to the WWW as a multidimensional (most commonly two-dimensional) landscape. Obtaining information in this landscape is expressed as traversing interconnected paths toward locations that contain information objects, such as user homepages and commercial catalog sites. Users say things such as, "I went to his homepage," and "I came back to where I saw that picture." Some of these information objects are talked about as two-dimensional and others, as three-dimensional; for instance, people say "in Yahoo!" which suggests a three-dimensional container, and "at AltaVista" which suggests a point on a two-dimensional plane.

In a follow up study, Matlock and Maglio (1997) asked experienced and inexperienced web users to judge the sensibility of sentences containing metaphorical language (specifically regarding motion) about obtaining information on the WWW. Using a scale of one to seven, participants rated the sensibility of sentences containing verbs of motion. For instance, "John went to a new web site today"; "Do you want to climb up to the UCSC home page?"; and "I waited for the information to come to me". Sentences in which the web user was viewed as an agent, actively moving along a horizontal path, were rated as significantly more sensible than those in which the web user moved up or down, and as significantly more sensible than those in which the web user was passive. These results suggest that both experienced and inexperienced participants have clear and consistent ideas about how motion does and does not occur on the WWW.

Though there are many ways in which people might talk about the WWW (see Benyon & Höök, 1997), the fact that they naturally talk about it using particular metaphors is no accident. As Lakoff and Johnson (1980) and others have argued, such language is motivated by metaphorical thought.

Metaphor and Thought

Prior to the seminal work of Lakoff and Johnson (1980), metaphor was generally seen as nothing more than a literary device. Lakoff and Johnson radically changed this misconception, offering compelling arguments to show that metaphor is an integral part of our thoughts and actions:

> Metaphor is typically viewed as a characteristic of language alone, a matter of words rather than thought or action... that metaphor is pervasive in everyday life, not just in language, but in thought and action. Our ordinary conceptual system, in terms of which we both think

> and act, is fundamentally metaphorical in nature. (Lakoff & Johnson, 1980, p.3)

Subsequent work in cognitive linguistics and psychology has continued to offer theoretical and empirical evidence to show that metaphor is ubiquitous and serves many functions relative to our conceptual experience (Lakoff, 1987; Sweetser, 1990; Turner, 1987). One of the functions of metaphor is that it helps us think about relatively abstract conceptual domains in terms of relatively concrete domains (Gibbs, 1994). For instance, spatial concepts are often helpful when reasoning about time (Gentner & Imai,1992). On the standard view of metaphor, a relatively concrete source domain maps onto a relatively abstract target domain. Consider the often-cited metaphor THEORIES ARE BUILDINGS. In this metaphor, elements of the conceptual structure of BUILDINGS (source domain) map onto THEORIES (target domain). Linguistic evidence to support the existence of this metaphor includes statements such as: "You need empirical evidence to buttress your arguments", "The foundation of the theory is shaky", "His entire theory was toppled by the claim that Basque is a language isolate", or "Construct a different argument to support your theory". It makes sense that this mapping progresses from BUILDINGS to THEORIES because buildings are common in our everyday experience. In Western culture, buildings serve an important function: namely, we live and work in buildings. In addition, buildings offer protection from adverse effects of nature, and so on. Theories, by contrast, are important in the academic or philosophical world, but not commonplace to most people.

Another example of a metaphor is THE MIND IS A CONTAINER. In this case, the concrete conceptual domain of CONTAINER maps onto the more abstract conceptual domain of the MIND. Hence, we understand the mind as storehouse. Ideas can enter the storehouse, can be processed there, stored in a specific location, or even misplaced. Linguistic evidence for this metaphor includes expressions such as "The thought suddenly came into my head", "It's in the back of my mind", or "She lost her senses". This metaphor underlies many standard psychological theories (see Gibbs, 1994, for discussion).

As pointed out by Coulson (1996), the standard approach to metaphor arose in part to account for simple examples of analogical thinking, such as TIME IS SPACE. As such, the approach is parsimonious but cannot account for complex mappings requiring some degree of sensitivity (Turner & Fauconnier, 1995). Moreover, the standard approach falls short with respect to productivity: Why do only certain elements of the source domain map onto the target domain? Consider THEORIES ARE BUILDINGS. As noted, foundation and support map onto the target domain, but doors and windows do not. Recent approaches have attempted to solve this problem by suggesting that there are a variety of types of metaphors, including primitive and compound (e.g., Grady, Taub & Morgan, 1996).

In any event, the more recent approaches to metaphor only diverge from the standard model with respect to issues of mapping complexity; all agree that metaphor helps structure how people think. But metaphor is only part of the story.

Image Schemata

Image schemata, basic pre-conceptual structures that arise from our embodied experience, shape both metaphorical and non-metaphorical thought (Gibbs & Colston, 1995; Johnson, 1987, 1992; Lakoff, 1987), and are formed early in development (Mandler, 1992).

Daily life includes active physical motion towards objects or destinations (concrete or abstract): going to the door to let the cat out, walking or driving to work, and reaching out to grab a pencil or pick up the telephone. Life also includes abstract motion toward goals (abstract destinations): working to get a promotion, writing a dissertation to obtain a degree, and saving money for a trip. Each of these actions involves the image schema TRAJECTORY, comprised of a starting point, an end point, and a path between the two. Another image schema is CONTAINER, which arises out of bodily experience: swallowing things, entering and remaining in buildings, and so on. As we will see later, these image schemata figure prominently in how people view obtaining information on the web.

In what follows, we explore the nature of people's metaphorical conception of the WWW. We set out to investigate how users with varying levels of expertise talk about the web. We first describe a study involving verbatim reports elicited from both experienced and inexperienced web users when relating what they did while using the web. We next discuss reasons people use the metaphors they do, and finally, some implications of our results for the design of tools for navigating information spaces.

Study: Language Use on the Web

The purpose of this study was to discover more about how people think about the WWW in natural settings. We looked specifically at how people conceive of the actions that they take while using the web; for instance, to what extent they see themselves actively moving through space and to what extent they focus on the physical environment. We also wanted to observe differences between beginning and experienced users.

We hypothesized that beginners would talk about their experiences using the web in terms of the physical actions they performed more than experienced users would because beginners are likely to have only a partial understanding of the web domain. Along the same lines, we hypothesized that experienced users would generate more metaphorically consistent utterances than beginners would.

We analyzed the data both quantitatively and qualitatively. In the quantitative analysis, we counted utterances of various types to compare beginning and experienced web users. In our qualitative analysis, we followed a method similar to that of Raubal, Egenhofer, Pfoser and Tryfona (1997), who analyzed the image schematic structure of talk about wayfinding in airports.

Method

Participants used the WWW and were then asked to tell us what they did. The resulting discourse was analyzed primarily in terms of image schemata.

Participants

Twenty-four undergraduates at the University of California at Santa Cruz took part in this study, including thirteen males and eleven females. All were native speakers except five early bilinguals who demonstrated English proficiency equal to that of native speakers.

Procedure

Participants first completed a questionnaire on computer and WWW use. Questions included length of time using the web (e.g., one month or less) and hours per week generally used. Participants then sat at a computer that was running the Netscape Navigator browser, which displayed the homepage for the University of California, Santa Cruz. They were instructed to click on whatever icons or hyperlinks appeared interesting and to continue doing so for five minutes. The experimenter was extremely careful to avoid language that would bias the participant to think of the web metaphorically, such as, “go to that page”.

After the participant had spent sufficient time getting used to the task and experiencing the environment, he or she was instructed to look at a new domain: Yahoo!, a well-known catalog in which information is organized hierarchically. The participant was again instructed to use the mouse to gain access to information that seemed interesting and to continue to do so for five minutes.

A tape-recorded interview followed the WWW session. To begin, the experimenter prompted the participant: “Tell me what you just did using as much detail as possible.” If a response was not immediately forthcoming, the experimenter said, “Tell me what you did first.”

Coding

In coding the data, we distinguished among seven kinds of verbs and verb-preposition combinations that indicate seven kinds of web actions (see Table 1). Only utterances that referred to what the participant did while using the web were assigned to one or more of these categories of action types. For example, clauses such as “I'm kind of on a tight budget”, or “It's pretty easy” were not considered.

Category	Example Verbs
Outside	click, press, type, scroll, sit
TRAJECTORY	go, come, bring, follow
User Agent	go, follow
Web Agent	bring, come up, bring, show
CONTAINER	have, contain
Information Action	look for, lookup, search
Miscellaneous	look, see

Table 1: Verb coding scheme.

First, we distinguished between verbs that refer to *outside actions* and those that specify *inside actions*. Responses referring to typing on the keyboard, using the mouse, and clicking on browser icons constituted outside actions (e.g., “I *typed* something”, “I *clicked on* the grapes icon”, or “I *pressed* buttons”). Responses referring to a TRAJECTORY in web space were coded as inside actions (e.g., “I *went into* this thing called Yahoo”, “I couldn't *get back to* where I was”, or “It *brought me to* the Anthropology page”). In addition, expressions conveying information movement were also coded as TRAJECTORY

(e.g., "it told me", "it said"). TRAJECTORY responses were further split in two: those in which the user is the agent (e.g., "I *went*…"), and those in which the web is the agent (e.g., "It *took me to*…").

The fifth type of verb refers to a web site as a CONTAINER (e.g., "Yahoo *had* what I wanted"). The sixth refers to information actions (e.g., "I *looked up* Chewbacca"). The final category coded those verbs not categorized in the previous six. This miscellaneous category mainly contained clauses beginning with the expressions "I *saw* ..." or "I *looked at*..." because it is unclear whether these refer to visual perception of the screen (an outside action) or to visual perception of objects in web space.

Finally, note that we might also have used prepositions to help code for TRAJECTORY and CONTAINER. For instance, *through* and *to* suggest a TRAJECTORY, and *in* suggests a CONTAINER. For simplicity, we chose to rely solely on verbs. A complete analysis would undoubtedly include prepositions (and nouns) as well (see Raubal, Egenhofer, Pfoser & Tryfona, 1997).

Results

Participants were separated into two groups according to web experience: twelve beginners reported six months or less of web use, and twelve experienced users reported more than six months of web use.

Qualitative Results

To get a feel for the data and our coding scheme, consider the following utterance typical of beginners (Participant 4):

> … I clicked on uh grapes … and it brought me to um … this place where they had choices and then I clicked on bookstore…

Note two outside actions ("click"), one TRAJECTORY in which the web is the agent ("it brought me to"), and one CONTAINER ("they had"). In this utterance, the user clicks on an icon the screen, is taken to a new location, and then clicks again.

Now consider the following utterance produced by an experienced web user (Participant 14):

> … I went to net search because that seemed like a good wholesome opportunity for going somewhere else… I probably typed something and it told me I couldn't do it, so I dunno, I just went and clicked around a whole bunch…

In this case, there are three outside actions ("typed", "clicked", and "do"), two TRAJECTORYs with the user as agent ("went", "going") and one TRAJECTORY with the web as agent ("told"). Note that the verb "seem" and the second instance of "went" were not coded, as these do not refer to actions taken when using the web. In this utterance, the user goes somewhere, types something, receives information, and clicks.

In both cases, outside action is mixed with action inside the web's information domain. For the beginner, the web is a kind of conveyance that moves the user ("brought me to"). For the expert, the web is a kind of roadway on which the user moves ("I went"). For the beginner, the web passively contains information ("had choices"). For the expert, the web actively provides information ("it told me").

Consider another beginner (Participant 2):

> …I went into the um Brian's tattoo something or other, but when I clicked into it, it said that like it was gonna show tattoos of his body and like front, side, whatever… it had objects to click on, and I clicked on em and there was no pictures…

There is one TRAJECTORY in which the user is the agent ("went into"), and two in which the web is the agent ("it said", "was gonna show"). There were two outside actions ("click"), and one CONTAINER ("had"). As in both previous cases, outside actions are mixed with inside actions. Like the first beginner, this one views a web site as a container. Unlike the first beginner, however, this one views the web as a kind of roadway ("went") rather than as a kind of conveyance ("brought me to"). For this beginner, as for the expert, the web actively provides information ("it said").

The utterance from Participant 2 illustrates something our coding scheme does not recognize: the construction "click into". Whereas the verb "click" refers to an outside action, the preposition "into" refers to an inside location. Usually the verb "click" is followed by the preposition "on", and the construction refers to an icon or hyperlink visible on the screen. In this case, however, "click into" refers both to something visible on the screen and also to something contained in the

information space of the web. We will return to this point in the discussion of conceptual blends.

Finally, consider a second expert (Participant 23):

> …I couldn't get through. I returned to the first page I started on and selected travel.

In this case, there are just two TRAJECTORYs in which the user is the agent ("get through", "returned") and one outside action ("selected"). The path is blocked ("couldn't get through"), and previous steps were retraced ("returned").

In summary, both beginners and experts use the same sort of language overall. Most participants mixed outside actions with TRAJECTORYs, and CONTAINERs inside web space. In talking about the web, the web can move the user, or the user can move on the web; the web can simply contain information, or it can actively convey information. In any event, people seem to prefer to talk about their experience in using the web in more familiar terms, such as physical motion, physical actions, and physical containers.

Quantitative Results

The total number of verbs in each category was computed for each group (see Table2).

	Beginners	Experts
Outside	54	26
TRAJECTORY	56	87
User Agent	37	79
Web Agent	19	8
CONTAINER	22	11
Info Action	30	42
Miscellaneous	24	20
Total	186	186

Table 2: Total number of verbs in each category.

Because we collected frequency data, χ^2 was used to compare beginners and experts along each of the seven dimensions (see Table 3). As shown, differences were found for TRAJECTORY versus outside actions, user agent versus web agent, and CONTAINER versus all other verbs. Experts used the TRAJECTORY schema rather than outside action verbs more often than beginners. Within the TRAJECTORY schema, experts viewed themselves as the agent (i.e., moving through information space) rather than the web as the agent (i.e., the information moves through the web) more often than beginners. By contrast beginners more often viewed web sites as CONTAINERs than experts.

	Beginners	Experts	χ^2
TRAJECTORY vs. Outside	51%	77%	16.49**
User agent vs. Web agent	66%	91%	13.60**
CONTAINER vs. all others	12%	6%	4.02*
Info actions vs. all others	16%	23%	2.48
Miscellaneous vs. all others	13%	11%	0.41

Note: * $p < .05$; ** $p < .005$

Table3: Percentage of verbs in each category for each group. The χ^2 statistic compares the difference between groups.

Overall, web users—novice and experienced users alike—talked about using the web as if they had been moving from place to place though in fact they had not gone anywhere. The data also revealed a striking distinction between experienced and beginning users. Beginners more often mixed in their experiences using the keyboard, mouse, and other elements of the physical (non-web) domain (e.g., "I clicked on..." or "I typed in..."), whereas experienced users did not. In addition, beginners were more likely to refer to the web as a container than were experienced web users.

Discussion

Both beginning and experienced WWW users refer to WWW use in terms of physical motion, though experienced users produce more metaphorically consistent utterances. In what follows, we discuss some reasons people use the metaphors they use when talking about the web, as well as what this means for the design of tools for information navigation.

Agency and Web Use

According to our data, web users—even those who had never used the web before—often view themselves as moving along paths toward information. While traversing these paths, motion most often occurs on the part of the web user (even for beginners; see Table 3). Less often is the user viewed as the passive recipient of information or as a passenger in some web vehicle. This suggests that the web lacks agency.

One reason the user might view obtaining information as moving through space toward objects is the ease of getting around on the web. The most common way of moving from one web page to another is by clicking on hyperlinks or using the browser's back button (Catledge & Pitkow, 1995; Tauscher & Greenberg, 1997). Much less often do web users type in full addresses to obtain information. Clicking on links and seeing new information instantaneously creates a sense of fluidity and hence might create a sense of motion. This hypothesis can be tested by systematically varying the delay between clicking on a link and the subsequent presentation of information. The test would be whether longer delays result in fewer utterances in which the user is the agent.

Perhaps a deeper reason lies in embodiment (Johnson, 1987, 1992). People's web actions may reflect how people normally obtain things or move about in the world. Much of the time, they must stand up, walk towards, reach out, and grasp what they need or want. Such recurrent patterns of activity form part of everyone's mental experience in the form of image schemata, as discussed previously. Directed motion toward goals is part of our embodied experience, and as a result, it naturally plays a role in how we structure thought about the abstract domains, such as the web.

Conceptual Blends in Information Space

We now return to our finding that novice web users mix talk about the outside domain with talk about the inside domain more than experienced users did (see Table 3). Recall the utterances of Participants 2 and 4. These and all inexperienced web users often mixed inside and outside actions, seemingly unaware of the fact that hey were code-switching between them. Sometimes this sort of blending happened at the sentence level, as in "I clicked on (outside) grapes ... and it brought me to (inside)...". At other times, it occurred at the phrase level, as in "I clicked into it", in which the participant created a novel verb-particle construction. These findings reveal that in using the web, people naturally integrate two or more domains to create something more than simply the combination of its parts.

Such conceptual integration is not unique to web activity, or even to language use. For instance, Fauconnier (1997) shows how the theory of conceptual blends (Fauconnier & Turner, 1994, 1996) can account for the complexity of the familiar computer *desktop* metaphor. This metaphor is constructed on the basis of two separate conceptual inputs: (a) traditional computer commands, such as saving a file, and listing a directory; and (b) work in an office, including a desk, files, folders, and trashcan. How can these inputs be integrated and construed in a meaningful way? Fauconnier argues that a cross-mapping occurs such that computer files are mapped to paper files, directories are mapped to folders, and so on. General knowledge—such as image schematic notions of CONTAINER and TRAJECTORY—mediate the mapping. Structure is selectively projected from the inputs, yielding a coherent, well-integrated, emergent structure specific to the blend. This creates a perceived world in which a trashcan can sit on the desktop, in which double clicking opens files or applications, and in which objects are routinely dragged from one location to another. The integration is completely novel, but at the same time compatible and accessible to the desktop interface user. Note that if the mapping from the office domain to the computer domain were simple (i.e., creating no new structure), the computer desktop could be no better than a real desktop: such an interface could only selectively mirror the world.

We believe that the conceptual blend theory provides a nice tool for analyzing how web users think about the web. It explains how novice users frequently integrate inside and outside actions (e.g., "click into"). It might also explain why experts talk about the web less in terms of outside actions than novice web users: because experts rely on the input from the abstract web domain to a greater extent than they rely on input from the physical browser domain.

Conceptual blend theory also integrates web users' conceptual information much more effectively than would a standard metaphorical account, which would be limited to a single source domain and a single target domain (e.g., Lakoff & Johnson, 1980). Of course, such an account can explain the obvious metaphor, OBTAINING INFORMATION IS MOVING THROUGH SPACE, which refers to how users view themselves moving along paths to information objects. But it can say nothing about how web users naturally blend inside and outside actions.

Designing for Information Navigation

If metaphorical language in fact reflects metaphorical thought, and people naturally think of the web as a kind of physical space, what are

the consequences for the design of information navigation tools?

Shum (1990) points out many potential uses for the concepts of physical space in the structuring and presentation of information, such as Euclidean distance in two or three dimensions, direction, orientation, and depth. Nevertheless, Shum also notes that the key in adapting spatial metaphors to information presentation lies in understanding user tasks. Thus, adding a notion of distance to the information interface solely because physical space has distance would probably not be useful in all cases. For instance, distance in information space might reasonably be used to convey semantic relatedness (e.g., Chalmers & Chitson, 1992) or expected download delay (e.g., Barrett, Maglio & Kellem, 1997).

Our data show that even novice web users conceive of themselves as actively moving on the web under their own steam. Thus, we believe that the power of spatial metaphors for information presentation is not merely the result of people's *ability to use spatial metaphors*. Rather, its power lies in the fact that people *naturally use spatial metaphors*—that they cannot help but use them. It follows that interface designers should not construct virtual worlds that are merely consistent with ordinary experience and that merely use spatial attributes in task-relevant ways. Rather, the most useful information interfaces will target people's natural spatial understanding of information use and at the same time allow people flexibility to create an appropriate metaphorical understanding of the domain (see Kuhn, 1993).

For example, Dieberger's (1997) city metaphor for information navigation seems to follow the right approach. In particular, Dieberger carefully balances spatially real interface elements with *magic features* that break the spatial metaphor. In a sense, magic features provide the user with known boundaries that can be used in guiding the conceptual blending process. For instance, because magic windows provide shortcuts between distal points in the information city, semantic-relatedness need not be determined solely by spatial proximity. Nevertheless, both sorts of connections can be understood spatially as TRAJECTORYs, providing a consistent basis for the mappings.

Conclusion

This paper explored metaphorical conceptions of the WWW, as gathered primarily from language data. We showed that both experienced and inexperienced web users naturally talk about the web as if they are moving toward information, rather than as if information is moving toward them. We also found some differences between experienced and inexperienced web users in the consistency with which they talked about their web actions. These data were explained in terms of the same basic image schemata that structure thought in physical domains. In any event, our findings are relevant to the design of information navigation tools. If people naturally conceive of the web as physical space, tools for navigation can go much further than they do today in exploiting this connection. The trick lies in discovering the conceptual differences between real space and cyberspace, and then in using those differences to make the boundary apparent.

Acknowledgments

Thanks to Gilles Fauconnier, Ray Gibbs, and Mark Turner for thoughtful discussions on this topic, to Chris Dryer for advice on statistical analyses, and to the workshop program committee for many helpful comments on an early draft.

References

Anderson, J. R. (1980). *Cognitive psychology and its implications*. San Francisco, CA: Freeman.

Barrett, R., Maglio, P. P., & Kellem, D. C. (1997). How to personalize the web. *Proceedings of Human Factors in Computing Systems, CHI '97*. New York: ACM Press.

Benyon, D., & Höök, K. (1997). Navigation in information spaces: Supporting the individual. In *Human computer interaction: INTERACT '97*. Chapman & Hall.

Catledge, L. & Pitkow, J. (1995). Characterizing browsing in the World Wide Web. In *Proceedings of the Third International World Wide Web Conference.*

Chalmers, M. & Chitson, P. (1992). Bead: Explorations in information visualization. In *Proceedings of the Fifteenth Annual ACM SIGIR Conference on Research and Development in Information Retrieval*. New York: ACM Press.

Coulclelis, H., Golledge, G., Gale, N., & Tobler, W. (1987). Exploring the anchor-point hypothesis of spatial cognition. *Journal of Environmental Psychology, 7*, 99-122.

Coulson, S. (1996). The Menendez brothers virus. In A. Goldberg (Ed.) *Conceptual structure, discourse, and language*. Cambridge England: Cambridge University Press.

Dieberger, A. (1997). A city metaphor to support navigation in complex information spaces. In S. C. Hirtle & A. U. Frank (Eds.) *Spatial information theory: A theoretical basis for GIS (COSIT '97)*. Berlin: Springer-Verlag.

Fauconnier, G. (1997). *Mappings in thought and language*. Cambridge England: Cambridge University Press

Fauconnier, G. & Turner, M. (1994). Conceptual projection and middle spaces. Technical Report 9401, University of California, San Diego, Department of Cognitive Science.

Fauconnier, G. & Turner, M. (1996) Blending as a central process in grammar. In A. Goldberg (Ed.) *Conceptual structure, discourse, and language*. Cambridge England: Cambridge University Press.

Gentner, D. & Imai, M. (1992). Is the future always ahead? Evidence for system-mappings in understanding space-time metaphors. In *Proceedings of the Fourteenth Annual Meeting of the Cognitive Science Society*: Hillsdale, NJ: LEA.

Gibbs, R. W. (1994). *The poetics of mind*. Cambridge, England: Cambridge University Press.

Gibbs, R. W., & Colston, H. L. (1995). The cognitive psychological reality of image schemas and their transformations. *Cognitive Linguistics,6*, 347-378.

Grady, J., Taub, S., & Morgan, P. (1996). Primitive and compound metaphors. In A. Goldberg (Ed.) *Conceptual structure, discourse, and language*. Cambridge, England: Cambridge University Press.

Johnson, M. (1987). *The body in the mind: The bodily basis of meaning imagination and reason*. Chicago, IL: University of Chicago Press.

Johnson, M. (1992). Philosophical implications of cognitive semantics. *Cognitive Linguistics, 3*, 345-366.

Kuhn, W. (1993). Metaphors create theories for users. In A. U. Frank & I. Campari (Eds.) *Spatial information theory: A theoretical basis for GIS (COSIT '93)*. Berlin: Springer-Verlag.

Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago, IL: University of Chicago Press.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago, IL: University of Chicago Press.

Maglio, P. P., & Barrett, R. (1997a). On the trail of information searchers. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*. Mahwah, NJ: LEA.

Maglio, P. P. & Barrett, R. (1997b). How build modeling agents to support web searchers. In *Proceedings of the Sixth International Conference on User Modeling*. New York: Springer Wien.

Mandler, J. M. (1992) How to build a baby: II. Conceptual primitives. *Psychological Review, 99*, 587-604.

Matlock, T., & Maglio, P. P. (1996). Apparent motion on the World Wide Web. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*. Mahwah, NJ: LEA.

Matlock, T. & Maglio, P. P. (1997). Untangling talk about the World-Wide Web. Unpublished manuscript. University of California, Santa Cruz, Psychology Department.

Raubal, M., Egenhofer, M. J., Pfoser, D., & Tryfona, N. (1997). Structuring space with image schemata: Wayfinding in airports as a case study. In S. C. Hirtle & A. U. Frank (Eds.) *Spatial information theory: A theoretical basis for GIS (COSIT '97)*. Berlin: Springer-Verlag.

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

Sweetser, E. (1990). *From etymology to pragmatics: Metaphorical and cultural aspects of semantic structure*. Cambridge England: Cambridge University Press.

Tauscher, L. & Greenberg, S. (1997). Revisitation patterns in World Wide Web navigation. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI '97)*. New York: ACM Press.

Turner, M. (1987). *Death is the mother of beauty: Mind, metaphor, criticism*. Chicago, IL: University of Chicago Press.

Turner, M. & Fauconnier, G. (1995). Conceptual integration and formal expression. *Metaphor and Symbolic Activity, 10*, 183-204.