CZWeb: a Web-based workspace for media-rich communication and decision evolution

Brian Fisher
Media and Graphics
Interdisciplinary Centre
University of British Columbia
Vancouver, BC, V6T 1Z4
fisher@cs.ubc.ca

John Dill
School of Engineering Science
Simon Fraser University
Burnaby, BC, V5A 1S6
dill@cs.sfu.ca

EXECUTIVE OVERVIEW (150 WDS)

Despite the ubiquity of the desktop metaphor, the World-Wide Web has resulted in increased interest in browser-based interaction with information. Current browser designs have limited navigational and visualization support, and users tend to passively follow links through an ill-understood Webspace.

To actively interact with Web-based information, we need a new set of interface concepts. The less constrained tasks and changing information on the Web entails a more flexible environment in which users interact with a varying number of chunks of information around an area of interest: reshaping the space, clustering information, and annotating the new patterns. These "WebDesks" will grow and shrink over time, but it is important that users maintain their sense of place.

The fluidity of our interface requires us to rely less on visual semantics for design methodologies than upon more basic aspects of human cognitive architecture: spatial cognition, visual attention and indexical cognition (FINST) models.

TEAM INFORMATION:

Major current team members are Dr. John Dill and Dr. Brian Fisher. Their backgrounds are in graphics/visualization and Cognitive Science/HCI respectively. Short Bios follow. Others who contributed significantly include Dr. C. Jones, Dr. G. Collaud and Mr. P. Tan. Dr. Jones is a Senior Software Engineer Chesapeake Decision Sciences with a background in algorithm animation; Prof. Collaud is on the faculty of the University of Fribourg in Switzerland; Mr. Tan was a graduate student at Simon Fraser University during the initial development.

John Dill is a Professor in the School of Engineering Science at Simon Fraser University. Prior positions include MPR Teltech, Cornell University and several years at the General Motors Research Laboratories where he was involved in computer graphics research related to Computer-Aided Design. His current research interests include information visualization, engineering visualization, user interfaces and intelligence in computer-aided design. He is active in ACM's SIGGRAPH and is a member of the IEEE Computer Graphics and Applications Editorial Board. Dill received a B.A.Sc. in Engineering Physics from UBC, a M.S. from North Carolina State University and a Ph.D. in Engineering Science from Caltech.

Brian Fisher is the Associate Director of MAGIC, the University of British Columbia's Media and Graphics Interdisciplinary Centre. Prior positions include the Institute for Human Factors and Interface Technology at Simon Fraser University, Rutgers Centre for Cognitive Science, and the Centre for Cognitive Science at the University of Western Ontario. He holds a Ph.D. in Experimental Psychology from the University of California at Santa Cruz.

LEVEL OF INNOVATION:

Traditionally, interface designers have had as their goal the creation of software "tools' that would allow users to perform particular tasks easily and effectively. Methods and theories from Cognitive Science and Psychology that deal with the conscious performance of mental tasks (e.g. GOMS and keystroke models, task analysis, protocol analysis) and visuomotor tasks (e.g. Fitts Law, SR Compatibility) were adapted to create core methodologies in the study of human-computer interaction [1].

To date, the adaptation of theories and methods from Cognitive Science has been highly selective. The emphasis on semantic similarities (e.g. of a word processor with a typewriter, or a visual calculator with a paper spreadsheet) has reduced the need for interfaces based upon characteristics of human attention and perception.

Thus, perceptual cognition research has played a lesser, largely supportive role. Mapping the functions of the software onto real-world tasks is made easier when the stylized icons are perceptually similar to real-world objects and dissimilar to each other. Findings from perceptual research was employed to help designers to accomplish this. However the way in which these icons are interpreted by the human perceptual system is in no way central to the process of choosing and implementing the visual metaphor. Instead, they typically take the form of general guidelines [1], such as limiting short-term memory load to less than seven items, maintaining SR compatibility, and limiting the use of attention-demanding stimuli such as flashing icons. Indeed, some respected HCI researchers have even suggested that deeper theories from Cognitive Psychology are of little use in software design [2].

Challenges of the Web

Designers of an interface that complements the fluid and ambiguous nature of Web-based materials would find it difficult to create the sorts of semantic analogues that characterize the desktop metaphor. Unlike data and applications on a personal computer, Web information:

- May change in an unpredictable way, and at any time
- May move to a new location on the Web
- May be distributed in a variety of formats (e.g. JPEG, GIF, TIFF) that are logically equivalent to the user
- May be actively generated by an interactive application or consist of static html
- May be generated locally (e.g. Java) or server-side (e.g. CGI). This may or may not be important to the user.

Mapping these characteristics onto a static interface such as a desktop would create expectations of constancy that would not be kept.

Accordingly, we set out to create a fluid interface environment, an “anti-desktop” if you will, that was nevertheless understandable and navigable. Despite changes in the size, positions, and characteristics of objects in the interface, users must be provided the perceptual support needed to function within the environment. This requires a new set of design parameters based in the cognitive architecture of human performance in complex environments.

Spatial cognition in fluid environments

In a complex environment, multiple perceptual objects and events must be individuated. Each object will have a position in space and may have a variety of visual, auditory, and tactile characteristics associated with it. These characteristics must be parsed from the environment and correctly assigned to the mental representation of the external event.that gave rise to them, despite changes in its location or other characteristics[4,5]. In order to unconsciously “keep track” of a given object’s identity and history, observers are thought to generate an individual representation in memory (called an “object file” [3]) for each important object in view that contains its history.

It seems that there is a significant amount of interaction between cognitive and perceptual processes that is required to function in a rich perceptual environment. Yet, the nature of the divisions in the flow of information and control that make up the human cognitive architecture requires that cognition have limited access to the sensory processes that create these representations[6,7]. In order to understand operator performance in rich sensory environments, analysis of users' conscious performance of mental operations associated with their tasks must somehow be integrated with an understanding of their perceptual and motor processes. This requires a critical mechanism for linking thought processes with specific events in the perceptual world.

The minimal mechanism for a deictic perceptual/cognitive system is a set of spatial indexes, pointers, or attentional tokens that serves to link mental operations (ranging from simple visual routines such as collinearity to more complex conceptual structures) and specific perceptual events. These pointers were first used in computer vision applications by Pylyshyn and colleagues [8, 9,10]. These attentional tokens were given the name FINST, for FINger of INSTantiation. They have subsequently been used to explain human perceptual and cognitive processing in a wide variety of domains [11].

According to the FINST hypothesis, there are a limited number of attentional tokens (FINSTs) that constitute a fundamental bottleneck in human processing. Just as short-term memory limitations (i.e. 7 +- 2 items [12]) and focal attention (i.e. a “spotlight” of attention) limit our ability to perform certain tasks, the number of spatial indices can be a determining factor in our ability to parse complex displays. Current thinking suggests that at most six items in a display can be simultaneous individuated by FINSTs for cognitive processing. These items receive preferential processing in a number of ways that will be described below.

According to this hypothesis, items in a complex display will fall into three separate categories:

- Attended items-- Items that occur in a spatially contiguous region that the operator is currently attending to. Processing in this region follows the traditional “spotlight” metaphor of endogenous attention, where the level of processing (grossly defined) varies inversely with the size of the area attended.

- Indexed items-- Up to six display items can be individuated and indexed by having a FINST assigned to them. FINSTs are usually assigned in a bottom-up manner based upon salient display events such as the onset of a new item. FINSTed items gain a number of specific processing advantages: They are available (potentially in parallel) as arguments for simple perceptual routines such as collinearity, conjunctive search, subitizing, and for rapid selection for focal attentive processing. These items generate unique mental representations called object files [3] that allow multiple characteristics (i.e. perceptual features) to be monitored in parallel and maintained through time (i.e. as an object-bound history in short-term memory). Unlike attended items, FINSTed items can be distributed across the screen in any configuration (i.e. they do not need to fall within a contiguous region to receive enhanced processing).

- Background items-- The remainder of the display receives very limited processing. New items that appear at unindexed locations may draw an index if their onset is salient; however if onset occurs during a saccade or a screen blink the item itself is unlikely to be noticed [13]. Changes in existing objects are also likely to go unnoticed, and will in all cases be responded to more slowly than similar changes in FINSTed items. Finally shifting attention to these items is slower than attending FINSTed items, and the time to required to attend to them will be roughly proportional to their distance from the current focus of attention [14].

In tests with human subjects, display items that are FINSTed have the potential to be accessed and acted upon by cognitive processes in parallel, with higher priority than unFINSTed items. We can use this model to predict the interaction of dynamic display events and cognitive processes, and to design displays that are optimized for the particular mental processes.

A concentration on the nature of the linkages between perceptual events and cognitive structures substantially alters the way we think about mental representations. Given this mechanism, we can derive models of information processing that rely upon the perceptual world to provide much of its own representation [15,16]. The theoretical issues involved are beyond the scope of this paper, but a forthcoming book by Pylyshyn (to be published by Ablex) should provide sufficient background for interested readers.

An interesting repercussion of the FINST hypothesis deals with the way in which we think about and remember physical spaces. Our introspections tell us that the way in which we think about and remember objects in space (e.g. the arrangement of furniture in our living room, or the path we take to work each day) is qualitatively different from the way we think about and remember facts or experiences we have had. This difference is born out by a great deal of research in spatial cognition [17].

The nature of that difference, however, may not be intuitively obvious. While we typically feel that our memory of events in space is itself spatial (e.g. we "visualize" our living room in a way that is very similar to actually seeing it), this intuition is inconsistent with the modularity hypothesis and the evidence that supports it. Instead, it is suggested that cognition is sensitive to qualitative spatial relationships between objects (e.g. "inside of", "above" beside") that are the products of simple "visual routines" [18].

Mental models of Web space

Suppose, for example, we intended to devise a graphical "map" of a complex data space (in this case the World Wide Web) taking into account what the FINST hypothesis says about how we parse complex displays and think about and remember physical spaces. How would this map differ from a static (paper) map or a physical space? Physical space has metric characteristics (i.e. the distances between items can be measured to any arbitrary level of accuracy).There is no obvious analog to metric distance in Web space, yet we do think of the Web a space of sorts.

According to the FINST hypothesis, the metric characteristics of events in physical space receive limited processing at the cognitive level. Processing spatial relationships takes place at the perceptual stage, and FINSTs serve to individuate a small number of salient events for cognitive processing. The alternative view would render the FINST mechanism unnecessary, as events could be individuated at the cognitive level by simply remembering their spatial coordinates.

At the cognitive level, we hypothesize that spatial relationships are preserved in a qualitative sense, but metric positions of events are not readily available. Individuation of objects and events may be spatially derived at the perceptual level, but are communicated to cognitive processes by way of pointers (FINSTs) that do not specifically pass on metric information.

Web users may indeed think of the Web as a spatial entity in the sense that individual Web pages can be individuated (i.e. they can be thought to be in different places), and exist in some relationship to other pages (e.g. they are on the same server), but without the metric characteristics of real space. Lacking a spatial display representation of the Web, they appear to be dealing with Web space at the cognitive level only.

If these theories are correct, they should help us to design representations of data spaces that are tailored to the way in which we process, think about, and remember events in complex scenes. We can extrapolate from these theories to generate some design recommendations:

First, the number of important display items would be kept small, to avoid exceeding the number of FINSTs available (approximately 6). this would insure that each display item could be indexed. This in turn would make it likely that each item could be tracked across eye movements or display changes, that it would maintain an individual object-history representation, and that changes in its display characteristics would be noticed quickly.

Second, since metric distances are relatively unimportant, changes in the display configuration could be allowed insofar as they do not interfere with the more qualitative spatial relationships (e.g "inside-of", "on-top-of") that are hypothesized to be preserved in memory. This allows for limited alterations in the size and layout of display items when circumstances warrant (for example, when more screen real-estate is needed).

Third, if continuity of processing is important, we must avoid display transformations that disrupt the ability of FINSTs to "stick" to their object: rapid onsets of irrelevant items should be avoided if possible, and indexed items should move smoothly to new positions rather than abruptly shifting location.

Patterns of use of the Web

Studies of Web use report patterns of behaviour that support the idea that users are often interested in assembling and revisiting a limited number of pages-- what Saul Greenberg calls a "recurrent system" [19]. According to research in U.Calgary's Grouplab [20]:

1. About 30% of all logged navigation activity involve use of the Back button
2. About 60% (s.d. = 9%) of all page visits are to pages a person has been to before
3. A list of 6 or so URLs just visited contain the majority of pages a person will visit next

4. Frequency of visits is a poor indicator: most people only access very few pages frequently and regularly.

5) Users do not exhibit strongly repeatable linear pattern of visits; while they tend to revisit recent pages, they do so in a different sequence

The recurrent system model of Web use postulates that users are primarily interested in creating sets of small numbers of familiar sites, bringing in new information from other sites as it is needed. If we consider the finding in #3 above as a rough estimate, we find that most of their referents could in principle fall within the number of FINSTed display items.

CZWeb: a flexible map of Web space

CZWeb (figure 2) is a tool for mapping portions of the World Wide Web that implements these recommendations. CZWeb runs as a companion to Netscape, automatically generating a flexible "map" of web sites as they visited. CZWeb currently runs on a PowerMac, and tracks Netscape actions by monitoring the Apple Events that it generates.

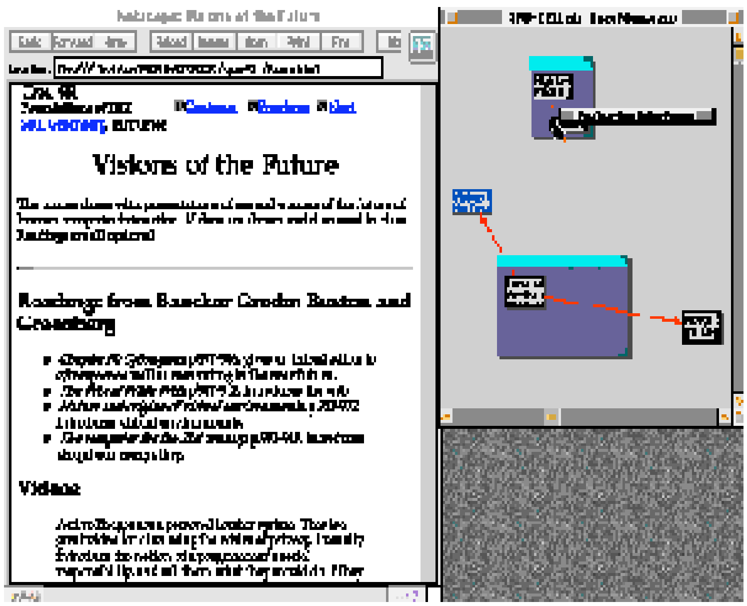

Figure 1: CZWeb and Netscape

In order to minimize the number of objects in the screen, CZweb uses a heirachial graph representation. Individual pages loaded into Netscape are represented as page nodes on the CZWeb display. One or more related page nodes can be grouped into cluster nodes (Figure 2).

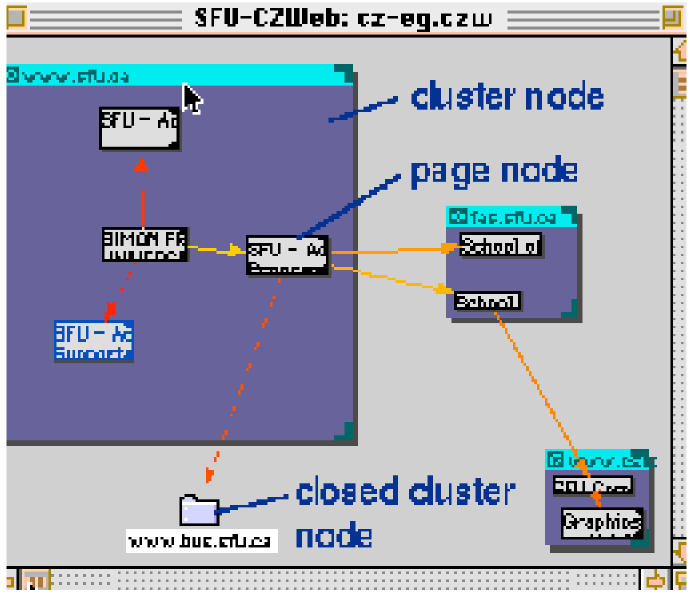

The default is to generate a cluster node for each Web server visited, and to display the pages served by that server within its cluster node. As new pages are added, CZWeb uses the Continuous Zoom display algorithm [21] to conserve real estate. It sequentially changes the position of existing pages, then reduces their size, and finally collapses smaller clusters into icons to free up real estate. These transformations are smooth, to maintain FINST binding, and preserve as much as possible the qualitative configuration of display elements (e.g on-top-of, beside, etc. Preserving the configuration of elements can help the user to cognitively "chunk" individual elements to further aid the recall process [22].

CZweb page nodes are active, and the user can reload a page in Netscape by clicking on its icon in CZWeb. The display is also configurable by the user-- they can move items, increase or reduce their size, close cluster nodes, create new clusters and populate them with pages dragged from other clusters. This enables them to create customized maps of important sites in their preferred organization and spatial layout. These maps can be saved to disk and mailed to other users.

User tests

User tests of CZWeb cannot in themselves confirm the effectiveness of the manipulations without appropriate control applications. However high user satisfaction ratings and observations of usage in laboratory contexts support our usability claims in a general sense. These findings were reported in more detail in a previous publication[23]. In summary:

- 12 university students acted as subjects
- Their task was to answer questions about the 96 Olympic sailing races
- Information came from a local Web site on the races
- Subjects were asked to think aloud and were videotaped as they solved the problem
- Subjects answered a pre-test and post-test questionnaire

The results of the questionnaire showed that users found CZWeb generally easy to use ($p < .05$), and useful for moving from site to site ($p<.00101$). Compared to Netscape navigation features (back/forward, history) Subjects rated CZWeb as more useful for helping to understand Web organization ($p<.0001$), moving to less recently visited sites($p<.0001$), and knowing where they are in the Web and avoiding going to the wrong site ($p<.0001$).

WebDesks as a novel interaction methodology

CZWeb was envisioned as a Web mapping and bookmarking tool that used spatial transformations to more effectively use limited screen real estate. The WebDesk concept extended the functionality of the interface to allow users to annotate page and cluster nodes, an idea suggested by Steven Forth of DNA Multimedia.

The effect of this small change in the interface is to allow users to link their own knowledge to bits of Web media. These hybrid representations are called KnowledgeObjects. With KnowledgeObjects, it becomes possible to interact with the WebDesk in a series of “passes” that iteratively alter the WebDesk to reflect the user’s changing understanding of the information. By re-annotating page nodes, moving page nodes to clusters, annotating clusters, creating new clusters and clusters of clusters, and moving them to new locations, users create visual analogues of their changing understanding/ interpretation of the material.

The effect of this is to create a “visual thinking space’ where changes in the relative location and clustering of items carries qualitative information, while more precise information can be carried in the annotations. User interaction with WebDesks are similar to those with concept mapping tools, with the exception that the KnowledgeObjects they contain fuse user concepts with chunks and clusters of media found on the Web.

Klein and colleagues [24] have argued that human cognitive abilities limit mental models to a maximum of three active agents (e.g. moving parts) and six discrete steps. This limitation constrains the effectiveness of formal inferential methods that are observed in laboratory tasks and taught in formal training programs. Klein’s research shows that skilled decision-makers rely instead on their trained perceptual abilities to support an understanding of the situation.

Traditional decision evolution environments are typically large (e.g. War Room) spaces where models of a situation are manipulated in a series of “what-if” analyses, downloading cognitive information to a physical representations [25]. Creating a similar “thinking space” on a small computer screen requires detail-in-context spatial transformations that maximize the effectiveness of limited screen “real estate”. Pylyshyn’s FINST theory of indexical perception provides us with the basis for transforming the display in such a way as to minimize the impact on the perceptual/cognitive linkages (i.e. FINSTs and object files) that support the decision-making process.

INTERACTIVE EXPERIENCE:

Spectators will observe the screen as a participant creates a CZWeb WebDesk by accessing Web pages on the browser, annotating them with comments, restructuring and modifying the spatial configuration and hierarchy of the nodes. They will attempt to return to a previous page finding that they can do so despite changes in the configuration of the display. Subsequent users can access their WebDesk, or our prepared WebDesks to judge the ease of learning of the transformed space.

APPEARANCE:

A videotape will follow by snail mail.

HISTORY:

The original CZWeb Web mapping application was developed by John Dill, Chris Jones, Gerald Collaud, and Paul Tan at Simon Fraser University. The core Continuous Zoom algorithm was developed by John Dill, Lyn Bartram, Ho, A., and Frank Henigman in 1994.

The application of new models of human cognitive architecture to interface design was the focus of a four-year Institute for Robotics and Intelligent Systems project, with the authors and Principal Investigators Kelly Booth, Jacquie Burkell, Tom Calvert, Roy Eagleson, and Zenon Pylyshyn.

The WebDesk metaphor evolved from discussion with the authors and Steven Forth, Hanif Jan Mohamed and Devon Girard of DNA Multimedia. ThoughtShare Inc., a DNA spin-off company, has licensed rights to the application for commercialization.

IMPACT:

The screen layout algorithms and inset zoom technology are themselves of interest, and the overall look and feel have generated positive responses in many users.

The application of the FINST visual cognition theory to display design is new, and may lead to other useful applications. We have begun work on several of these in collaboration with industrial partners ThoughtShare Inc. and Virtual Learning Environments Inc. We are also exploring the impact on the design of air traffic control systems with colleagues from Hughes Raytheon Labs.

WebDesks provide users with a single representation for diverse Web information, to better enable them to bring to bear their perceptual/cognitive abilities. Through the sequence of selecting, grouping, and annotating Web pages, changes in the user's understanding is continually supported by the configuration and nature of the information represented on the display.

As a communication medium, WebDesks allow users to reference specific materials on the Web directly, and their incorporation in online conference can support multimedia-based collaboration and training. Because discussions takes place "over top of the web sites", comments can be changed without consequence to the web site, and the contents of the site can be altered without damaging its association with the user's annotations.

Other uses for the WebDesk format might allow them to be placed on the server to be downloaded and modified by users. These may contain fixed icons such as landmarks, trademarks, etc. Search Engines can also serve results in the form of WebDesks, for reduced server load and increased usability.

TECHNICAL:

The display requires a PowerMac computer with Netscape Navigator 3 or better and Interned access. A display projector and screen would greatly enhance the presentation.

ACKNOWLEDGMENTS

This research was supported by Hughes Raytheon Labs, the Institute for Robotics and Intelligent Systems (Human Machine Interface-1 project) Network of Centres of Excellence, and the Telelearning Network of Centres of Excellence.

REFERENCES

1. Shneiderman, B. *Designing the User Interface* Addison Wesley, Reading, MA, 1998.

2. Landauer, T. K. Let's get real: A position paper on the role of Cognitive Psychology in the design of humanly useful and usable systems. in J.M. Carroll (Ed.) *Designing Interaction*. Cambridge, Cambridge MA 1991.

3. Treisman, A. and DeShepper, B. Object tokens, attention, and visual memory. In T. Inui and J. McClelland (Eds.) Attention and Performance XVI: Information Integration in Perception and Communication, Cambridge, MA: MIT Press, 199. pp.15-46.

4. Fisher, B.D. *Integration of visual and auditory information in the perception of speech events*. Doctoral Dissertation, University of California at Santa Cruz. 1991

5. Fisher, B.D. and Pylyshyn, Z.W. The cognitive architecture of bimodal event perception: A commentary and addendum to Radeau. *Current Psychology of Cognition* 13:1. pp. 92-96.

6. Pylyshyn, Z. W. *Computation and Cognition*. MIT Press, Cambridge MA 1984.

7. Pylyshyn, Z.W. The role of cognitive architecture in theories of cognition. In K. VanLehn (Ed.), *Architectures for Intelligence*. Hillsdale: Lawrence Erlbaum Associates Inc. 1991.

8. Pylyshyn, Z.W., Elcock, E.W., Marmor, M. and Sander, P. A system for perceptual-motor based reasoning. (Tech. Rep. no. 42) London Ontario: Dept. of Computer Science, University of Western Ontario.1978.

9. Pylyshyn, Z.W. The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. Cognition, 32, 1989 pp 65-97.

10. Pylyshyn, Z.W. "Here" and "There" in the visual field. In Z.W. Pylyshyn (Ed.), Computational Processes in Human Vision: An Interdisciplinary Perspective. Norwood, N.J.: Ablex Publishing. 1988.

11. Pylyshyn, Z., Burkell, J. Fisher, B. Sears, C. Schmidt, W. Trick, L.Multiple parallel access in visual attention. Canadian Journal of Experimental Psychology 48:2, 1993, pp.260-283

12. Miller, G. The magical number seven plus or minus two: Some limitations on our ability to process information *Psychological Science* 63, 1953, pp. 81-97.

13. Rensink, R.A., O'Regan, J.K., & Clark, J.J. Failure to Notice Scene Changes made during Image Flicker. Talk presented at Cambridge Basic Research. June 26, 1995 (abstract).

14. Tsal, Y. Movements of attention across the visual field. Journal of Experimental Psychology: Human Perception and Performance, 5, 1983 pp. 523-530.

15. Clark, A. *Being There: Putting Body, Brain, and World Together Again* MIT Press Cambridge, MA 1996

16. Suchman, L. *Plans and Situated Actions: The problem of Human-Machine Communication*. Cambridge, Cambridge MA 1987.

17. Pylyshyn, Zenon W. The imagery debate: Analogue media versus tacit knowledge. *Psychological Review* 88(1) 1981. pp. 16-45.

18. Ullman, S.Visual routines. *Cognition,* 18, 1984 pp. 97-159.

19. Greenberg, S. The Computer User As Toolsmith : The Use, Reuse and Organization of Computer-Based Tools (Cambridge Series on Human-Computer Interaction, No 6) 1993

20. Tauscher, L. and Greenberg, S. (1997). Revisitation Patterns in World Wide Web Navigation. In Proceedings of the ACM SIGCHI'97 Conference on Human Factors in Computing Systems, Atlanta, Georgia, March 22-27, ACM Press.

21. Dill, J. Bartram, L., Ho, A., and Henigman, F. A continuously variable zoom for navigating large hierarchical networks. *Proceedings of the 1994 Conference on Systems, Man and Cybernetics*, pages 386-390, 1994.

22. Ballard, D.H., Hayhoe, M.M., & Pook, P.K. Deictic codes for the embodiment of cognition. *Behavioural and Brain Sciences* (in press).

23. Fisher, B.D., Agelidis, M., Dill, J., Tan, P., Collaud, G., & Jones, C. CZWeb: Fish-eye views for visualizing the World Wide Web. In M.J. Smith, G. Salvendy & R. J. Koubek *Design of Computing Systems: Social and Ergonomic Considerations: Volume 2, Advances in Human Factors/Ergonomics,* 21B. pp 719-722, Elsevier. Amsterdam 1997

24. Klein, G. Sources of Power : How People Make Decisions. MIT Press Cambridge, MA 1998

25. Hutchins, E. *Cognition in the Wild*. MIT Press Cambridge, MA 1995