Index Keywords: cognitive principles, theory, guidelines, attention, memory, web page design

Cognition and the Web: Moving from Theory to Web Design

Mary P. Czerwinski

Kevin Larson

Microsoft Research

RUNNING HEAD: Cognition and the Web

Abstract

The links from basic cognitive theory to applied web design practice are not as strong as we would like to see. This chapter outlines some areas in which the two disciplines complement each other. We also describe several challenges such as designing websites from first principles and moving both fields forward in mutually beneficial ways. Although much progress has been made in the last few years, we believe that both disciplines can benefit from a closer partnership, and we hope that this chapter begins to detail some inspiration for that advancement. We will discuss many examples from our own research and identify when designing from first principles was successful, as well as when more work was needed to complete a good web design. We would like to encourage people performing basic research to use more challenging real-world tasks like web navigation and web search for their scientific inquiries. In addition, it is clear that theory and practice from many other disciplines are needed to bridge the gap, including the fields of computer science, anthropology, information visualization and design.

INTRODUCTION: THE WEAK CONNECTION BETWEEN COGNITIVE THEORY AND WEB DESIGN

Cognitive scientists who work in the web design industry typically find themselves designing and evaluating complex web-based systems to aid humans in a wide range of problem domains, like e-business, interpersonal communications, information access, remote meeting support, news reporting, or even gaming situations. In these domains, the technologies and the users' tasks are in a constant state of flux, evolution and co-evolution. Cognitive scientists working in human-computer interaction and web design may try to start from first principles developing these web-based systems, but they often encounter novel usage scenarios for which no guidance is available (Dumais & Czerwinski, 2001). For this reason, we believe that there is not as much application of theories, models, and specific findings from basic cognitive research to user interface (UI) design as one would hope. However, several analysis techniques and some guidelines generated from the literature are useful. In this paper we outline some efforts in human-computer interaction (HCI) research from our own industrial research experience, demonstrating at what points in the web design cycle HCI practitioners typically draw from the cognitive literature. We also point out areas where no ties exist. It is our goal to highlight opportunities for the two disciplines to work together to the mutual benefit of both. In addition, our argument holds true for several other research areas, including socio-cultural and technological fields' theoretical and research contributions. We argue primarily from a cognitive viewpoint as that is our background and because the domain is resplendent with low-hanging fruit that could be applied to web site design.

Throughout our discussion we provide examples from our own applied web research, as well as from the work of others. Our goal will be to highlight where the discipline of cognitive science aided our design efforts, as well as where more effort needs to be applied before a thorough set of web design practices and guidelines from any particular subfield of cognitive science can be instantiated.

PRACTICING WEB DESIGN WITH A STRONG COGNITIVE FOUNDATION

We cannot currently design complex web sites from first principles. Useful principles can be drawn from the sub-domains of sensation, perception, attention, memory, and decision-making to guide us on issues surrounding screen layout, information grouping, menu length, depth and breadth. Guidelines exist for how to use color, how to use animation and shading, or even what parameters influence immersion in web-based virtual worlds. Web designers also often borrow ideas from best practices, such as very successful web sites that have been iteratively refined and accepted broadly in the electronic marketplace. However, there will always be technology-usage scenarios for which the basic research simply does not exist to guide us during design. It is well publicized that users scan web pages, and most are not regular users of any one web site. These two facts alone make it fairly likely that undertaking good design for the web will be daunting compared to that of more frequently used software applications. In addition, web sites are used by a very wide range of users for a large variety of tasks—far more than the traditional office software application. Finally, the information itself that is on the web is an extremely complex material to work with, which makes the design task all the more challenging.

We feel that in order to move toward a stronger cognitive foundation in web design, changes in the way basic research in cognitive science is performed would have to be put into place. Basic cognitive science research has not even begun to scratch the surface in terms of complexity in their paradigms as is required when designing a web site today. In addition, the breadth of users and the individual differences therefore encountered are equally not addressed in most psychological research in information processing. The stimuli and tasks typically used in cognitive research are too simple, the mental load is often not high enough, and the user is not often driven to perform based on emotional, real world goals and constraints in the experimental scenarios used. It is for this reason that many of the theories coming out of cognitive psychology (and many other sub-disciplines in psychology) make predictions at far too low of a level in the knowledge worker's mental repertoire to be of much use to a web designer. Although the cognitive information processing approach has valuable findings and methods to pull from, the approaches, theories and findings will have to more closely approach realism to build a strong cognitive foundation in web design practice.

So while there are a large number of findings within cognitive science from which web designers can draw, knowledge falls far short of covering the entire practical needs of a web designer. Next we will review a few methods in the areas of cognitive science that have proven most useful to web user interface design.

EXPERIMENTAL METHODS

Many of the methods that cognitive scientists use during their academic careers are very useful to the web practitioners and researchers. This is often the first way in which a cognitive scientist in industry can add value to a product team. The inventors have a particular design idea for a web site, but they typically have little experience in designing the right experiments and tasks to validate those design intuitions. In our work we have used visual search tasks, dual tasks, reaction time and accuracy studies, deadline procedures, memory methods like the cued recall task and others to explore new web technologies and interaction techniques. Problems with the proposed web site designs are almost always identified, as rarely is a new web design flawless in its first stages. And, iterative design, test, and redesign works to improve the existing technology implementation from that point forward.

The downside of using traditional experimental designs and tasks in web work is that factorial designs with tight control and many subjects are simply not feasible given the time and resource constraints faced by design professionals. In addition, most real world designs consist of many variables and studying all possible combinations just isn't possible (or perhaps appropriate) given the short turn-around cycle web design teams are notorious for. Finally, it is often the unanticipated uses of a web site that are the most problematic, and these by their very nature are difficult to bring into the lab ahead of time. To understand some of the issues, think about how you would go about designing and evaluating a new voice-input web portal over a six-month period across multiple users, tasks, hardware systems and languages.

One might conjecture that important independent variables could be studied in isolation with the best of those being combined into a final design. In reality, this rarely works in web design. We have witnessed teams that have iteratively and carefully tested individual features of a web design until they were perfected only to see interactions and tradeoffs appear when all of the features were united in the final stages of design or when the site was used in the real world. Let's take a simple example of navigation bars. It may be that your labels and layout are intuitive and distinctive as evidenced by user studies of early prototypes, but when placed on the full web page, complete with banner advertisements, notifications, additional navigation content above, below or to the side of the original navigation bar, users cannot find what they are looking for. Navigation content is easily swallowed when housed in the framework of the attention-grabbing content of most web pages. It takes careful iterative test and design of the full web page, not just the single navigation bar, to get the visual layout of the bar to the level of priority necessary so that the user can complete their navigation tasks quickly and easily.

In addition to experimental methods, practitioners use a wide range of observational techniques and heuristics (e.g., field studies, contextual inquiry, heuristic evaluation, cognitive walkthrough, rapid prototyping, questionnaires, focus groups, personas and scenarios, competitive benchmarks tests, usage log collection, etc.) to better understand web site usage and to inform design. These techniques are often borrowed from anthropology or sociology, and rarely would a web design be complete without these research techniques and tools. Often a qualitative description of the jobs people are trying to do, or how a web site is used in the field by real people doing their real work is much more valuable than a quantitative lab study of one small web site feature. Observational techniques, task analysis, and related skills are much used in the practice of web design but little represented in most cognitive science curricula. This is an area in which researchers in cognition find their background and training lacking when first introduced to applied work. In addition, the methods themselves are not necessarily as well honed or useful as they could be in influencing web user interface design.

Evaluating and iterating on existing web user interfaces is only one aspect of the web design task. Generating a good design in the first place or generating alternatives given initial user experiences is equally important. In order to do this, web designers need to pull from basic principles in cognition, as well as basic principles of design!

There certainly are areas in which basic cognitive research has helped the domain of web design by providing specific findings and guidelines for design. Some of the best-known and often-cited examples of the applicability of results from basic psychology are the rather low-level perceptual-motor findings that have been used quite effectively in web design. For example, Fitt's Law and Hick's Law have been used to design and evaluate input device use during HCI for years. The Power Law of Practice and the known limits on auditory and visual perception have often been leveraged in the design of interactive web systems. Many other results are seen in guidelines for good layout, the use of color and highlighting, depth and breadth tradeoffs in menu design and hyperlink grouping, and many general abstractions have made their way into "libraries" of parameters describing typical response times, for example. Using findings and guidelines like these allow web designers to start with a good initial design, or prevent silly mistakes, but it doesn't guarantee a useful, usable web site when all the variables are combined together into one design.

Even though specific findings and guidelines have proven useful in some cases, there also exist many problems in their use, which limits their effectiveness. Guidelines are often misused or misinterpreted. For example, using the rule of thumb of having only 7+-2 items in a menu, or applying visual search guidelines but not taking into account the frequency with which items occur or the quality of category cohesiveness or labels. In addition, guidelines are often written at too abstract a level to help with specific designs, or alternatively they are too specific for a given usage context. A designer often finds it very difficult to combine all of the recommendations from specific findings or guidelines without knowledge of the costs, benefits or tradeoffs of doing so.

It is not our intent to state that there are no basic principles from which to design web sites available from theories in cognitive science; basic principles are available and do exist. We will highlight several of these principles below, noting when and where in the design cycle they were applied. Still, we argue that not enough of them exist, and that more work with complex user decision making, multi-tasking and mental load needs to be carried out, in addition to working with complex information stimuli. One specific area in which web designers have been able to draw quite readily from cognitive theory is in the domain of visual attention. Especially with regard to the manner in which users will typically scan a web page for content or banner ads, cognitive scientists have provided many great design principles.

For instance, Lim & Wogalter (2000) report two studies that looked at the placement of static “banners” on the web. Recognition was shown to be reliably higher for banners in the top left or bottom right corners of the displays. They then extended this finding by placing banners in 16 regions across the display—outer, intermediate and central regions and each region included 4 corners of banner placement. Results showed that recognition performance was reliably higher for banners centrally located over those in the outer regions of the display. The intermediate regions were not significantly different from the outer or inner regions. A separate analysis was carried out on the 4 corners for each region, and again the top left and bottom right corners were significantly better recognized than the other 2 corners. The results suggest that advertisers would do well to place their banner ads spatially on a continuum from top left, middle or bottom right of the display rather than in other areas of the layout. In addition, the authors argue that “visual pagers” or notifications could be made more salient by using this spatial location positioning. However, they warn that these notifications could also be very distracting in these locations, and so they caution this placement advice for tasks that rely upon the user giving their full attention to the primary task, for instance in monitoring tasks like air traffic control.

Visual attention theorists have also provided us with good examples for presenting web content in a way that will grab users’ attention without much effort --called preattentive processing by Treisman (1985). Typically, tasks that can be performed on large multi-element displays in less than 200 to 250msec are considered preattentive. Eye movements take at least 200msec to initiate, and random locations of the elements in the display ensure that attention cannot be pre-focused on any particular location, yet subjects report that these tasks can be completed with very little effort. This suggests that certain information in the display is processed in parallel by the low-level visual system. One very interesting result has been the discovery of a limited set of visual properties that are processed preattentively, without the need for focused attention. Visual properties that are processed preattentively can be used to highlight important web display characteristics. The kinds of features that can be processed in this manner include line orientations, lengths, size, number, terminators, intersections, closure, color flicker, direction of motion, etc. Experiments in cognitive science have utilized preattentive features to show that users can rapidly and accurately detect the presence or absence of a "target" with a unique visual feature within a field of distractors, can rapidly and accurately detect a texture boundary between two groups of items, where all of the elements in each group have a common visual property, and can count or estimate the number of elements in a display with a unique visual feature. Obviously you can't leverage all of these preattentive cues and features at once! And research is starting to show that some of these preattentive features take precedence over others. For instance, Callaghan (1990) has reported studies that suggest that brightness overrides hue information and that hue overrides shape during boundary detection tasks. However, knowing what features can work toward popping out various sections of your web display can give the designer an added advantage if warranted for a particularly important user task. To exemplify this, let's consider Figure 1, a snapshot of http://www.corbis.com. Note the way color is used to pop out the categories of interest, and that very little information is competing with your attention to this area of the display. In addition, the image selected for the home page has a background that blends in nicely with the chosen links for navigation, adding to the pop-out effect a sense of symmetry and grouping.

Figure 8.1. http://www.corbis.com, Corbis Corporation's home page.

Likewise, there is much we know from the early work of the Gestalt psychologists about good web design. The Gestalt theorists gave us the principles of proximity, similarity, common fate, good continuation and closure. Their "Laws of Organization" primarily focused on recognizing form from background. For instance, the principles of 1) area states that the smaller a closed area, the more it is seen as the figure; 2) proximity states that objects close to each other tend to be grouped together, and 3) closedness states that areas with closed contours tend to be seen as figure more than do those with open contours. The law of symmetry holds that the more symmetrical a closed region, the more it tends to be seen as figure, and the law of good continuation states that the arrangement of figure and ground tends to be seen which will make the fewest changes in straight or smoothly curving lines or contours. All of these principles can easily be put into place during good web design. By way of example, let's examine Figure 2, a screenshot of Fidelity's web site. This web site was iteratively redesigned after studies revealed that users were ignoring areas of the web site thought to be of core importance (Tullis, 1998). It was thought that layouts that "guided the eyes" to more important parts of the web site would generate more customers to key areas. Note the great use of grouping and symmetry in the web page's layout, resulting in 6 distinct groups (including the navigation bar at the top) balanced from left to right across the page. Note too how the new user is roped into getting started by a prominent "New to Fidelity?" link squarely centered on the page, in an area most likely to get scanned. The web site is leveraging visual perception and attention principles in a very strong way.

Figure 8.2. The Fidelity.com web site (www100.fidelity.com).

Coming out of the visual design field, Bertin (1983) proposed the “semiology of graphics” and systematically classified the use of visual elements to display data and relationships. His system consisted of 7 visual variables: position, form, orientation, color, texture, value and size. He joined these visual variables with visual semantics for linking data attributes to visual elements. Note the similarity between Bertin’s 7 visual variables and the preattentive features proposed by cognitive researchers. Clearly this is one place were web designers can benefit from principles converging from two completely separate disciplines.

Spatial layout can be successfully applied to web pages in order to effectively leverage human spatial memory, as has been demonstrated quite effectively over the last few years (e.g., Chen & Czerwinski, 1999; Czerwinski et al., 1998; Robertson et al., 2000; Tavanti & Lind, 2001). Using spatial layout of your web site for this purpose will ensure that users will develop a better conceptual model of your web site’s content over fewer return visits. Since a user’s willingness to return to a site hinges on its ease of use and navigability (Lohse & Spiller, 1998), spatial memory for the site structure may be key. Research by architects has shown that distinctive paths, nodes, landmarks, boundaries and districts are very important to leverage this remarkable human capability (e.g., Lynch, 1960; Passini, 1984). Vinson (2000), in particular, has laid out very clear guidelines for user interface design with good landmarks for navigation. He stresses the importance of using landmarks that are distinctive, fall on clear gridlines with proper axis alignment, scale appropriately for the world they are placed in, etc. In addition to the user of landmarks, we have found in our research that allowing the user to place content where they want it further enhances spatial memory, likely due to the further processing that is required (both motorically through manual placement on the page and cognitively in terms of deciding where the content should go). Of course there is a tradeoff here in terms of whether or not users will tolerate personalizing the web page or not, since this requires added effort. Still, if the user is bound to return multiple times, the effort may outweigh the cost of trying to navigate a site that is suboptimally designed for a particular user's task. We think the Austin MONKEYmedia web site could demonstrate an effective use of landmarks to enhance navigability of their website, except for the simple problem that the landmarks shift spatially as the user moves from subsection to subsection of the site. A snapshot of their homepage is shown in Figure 3. The web site uses a set of 5 clearly distinctive landmarks to help the user navigate through these cornerstones of the web site from any location. Interesting animation and audio effects underline the distinctive landmarks, although it is unclear if these additional cues are useful or usable.

Figure 8.3. The MONKEYmedia web site (the ultrasensory website available at http://www.monkey.com/FR4/main.htm).

Well known psychological principles do not always apply easily. Larson and Czerwinski (1998) provided an example analyzing use of the rule of thumb for provide no more than 7 +/- 2 links on a website. The 7+/-2 guideline germinates from a famous paper by George Miller (1956), and the "magic number" describes the number of "chunks" of information an average person can hold temporarily in short-term memory. We had observed that in several websites that we worked on that if the number of links at the top level of the site exceeded a small number that novice web users made more errors when trying to find particular content. We speculated that short-term memory was a factor in how many links can be presented on a page for optimal access to a document, and that the number of links should not exceed 7 +/- 2. To test this hypothesis we created three websites with 512 documents that were similar in content but the structure varied in the number of links between which users would have to choose. One website presented three levels of links with 8 links per level, a second presented 16 at the top level and 32 at the second level, and a third site had 32 at the top level and 16 at the second level. It turned out that users performed reliably slowest and were most lost when the number of links was held to 8. Users were fastest and least lost with the website with 16 top-level items, but was not reliably faster than the website with 32 top-level items. In George Miller's (1956) famous article about the limits of short-term memory, he reported a study where subjects were asked to try to match a new tone to a set of existing tones. If the set of existing tones is small, subjects are very proficient at the task but perform much worse as the set grows. Superficially picking between tones and picking between a list of links seems very similar. Our study suggested that these are very different tasks. When picking between web links subject's short-term memory does not appear to be an important factor. Instead, it is far more important to ensure that the labeling is good enough to tell people what information will be available if they follow the link. Indeed, recent work by Dumais et al. (Chen & Dumais, 2000; Dumais, Cutrell & Chen, 2001) reports that the good visual categorization (spatial grouping and labeling) of web content is key to optimizing the presentation of navigable search results. An example applying this idea (one of the designs that tested successfully in the series of studies carried out by Dumais et al.) is presented in Figure 4.

Figure 8.4. A successful design for presenting web search results (adapted from Dumais, Cutrell & Chen, 2001).

One new phenomenon recently identified by psychologists that directly applies to web design is what is referred to as "change blindness" (Hudson, 2001). Change blindness is a perceptual and cognitive phenomenon that occurs when movement that typically accompanies change is masked. The masking could be brought about by blinking or being momentarily distracted by something elsewhere in the visual field. Change blindness occurs when the user does not notice a difference to the visual information from before to after the attentional distraction. Psychologists have verified that this change blindness occurs both in the real world and in graphical displays, and that a visual disruption as brief at 80-240 msec can result in a user not noticing a change in the display (see the web site at http://www.acm.org/sigchi/bulletin/2001.6 for an interesting example of this phenomenon). What this implies for web designers is that if your users are traversing pages that are related to each other, and that only subtle differences need to be identified by the user, the download time between navigating from one page to another may result in the user not noticing that they've arrived at the new page! In addition, as pointed out by Hudson (2001), directing the user to problems during form validation is a ripe area for change blindness to be exhibited. The designer will need to come up with attention attracting visual (and maybe audio) cues to direct the user to those fields in the form that require the user's attention, else change blindness may result.

Another example from our work in managing interruptions illustrates some of these issues. The problem area of interruption while multitasking on the computer is one where much of what we have learned from basic psychological research has been applicable in website design. For example, 100 years of attention research has taught us about limited cognitive resources, and the costs of task switching and time-sharing even across multiple perceptual channels. But how does one go about designing an effective web user interface for alerts when the user may benefit from the information? And how does one design an interface that allows the user to easily get back to the primary web task after a disruptive notification?

In practice, specific findings guide our intuitions that flashing, moving, abrupt onset and loud audio heralds will attract attention. But how much attraction is too much? How does this change over time? Do users habituate to the alerts? How does the relevance of the incoming messages affect task switching and disruption? Some studies in the literature suggested relevance was influential, while others found that surface similarity was more important in terms of task influence. All of these studies used tasks that were not representative of typical computer tasks (i.e., they were too simple, demanded equal distribution of attention, etc.). We ran our own studies and showed the benefits of only displaying relevant notifications to the current task in order to mitigate deleterious effects (Czerwinski, Cutrell & Horvitz, 2000).

When, as part of the same web notification design, we attempted to design a “reminder” cue that was visual and could help the user reinstate the original task context after an incoming notification, specific findings were again not useful. We found through our own lab studies that using a “visual marker” as a spatial placeholder to get users back to a point in a primary web task was not enough. Instead, users needed a “cognitive marker” that provided more of the contextual, semantic cues related to a primary task (Cutrell, Czerwinski & Horvitz, 2001). No specific findings existed in the literature for exactly what visual anchors could be used in a web display to help users tap into their mental representations of a task, or how many visual retrieval cues would be needed to reinstate a web task context most efficiently. Again, prototypes needed to be designed and tested. Our studies over two years taught us many important principles in the design of notifications to computing systems. Most importantly, the system must learn some information about the user’s current context (e.g., how busy is the user? What applications are being used? What is the topic of focus? How important is the incoming information? How relevant is it to the current topic of focus or one that is coming up shortly? What is the best way to present the information to the user given his or her context? On what device? Etc.). Though there are obvious privacy concerns here, we have found that users are mostly enthusiastic about a system that learns their habits over time and respects their current level of "interruptability". We found in several studies that a good notification system should wait for a motoric or cognitive "break" before presenting information to the user (Czerwinski et al., 2000), and that information that is relevant to the user's current task will be easier to attend to and then dismiss in order to return to the primary task (Czerwinski et al., 2000). Studies also showed that users were more likely to forget their primary task content if interruptions or notifications came earlier into the task, as opposed to later (Cutrell, Czerwinski & Horvitz, 2001). We found that users will not be likely to trust a faulty intelligent system, so it is very important that an intelligent notification system has good default behavior from the onset (Tiernan, Czerwinski & Cutrell, 2001). Finally, our research showed that repeating information via notification that a user has either already seen or that is now irrelevant to ongoing tasks is significantly more frustrating than a first or timely delivery of the same information.

In both the design of web notifications, and the corresponding web reminders, paradigms from very basic cognitive research had to be utilized in new lab studies to examine the broader usage context and the specific user interface design chosen for the target task domains.

NEW WEB METRICS: SUBJECTIVE DURATION ASSESSMENT (SDA)

Due to the lack of fit between web design practice and findings in cognitive science, our basic research also attempts to develop new metrics for web interaction assessment. Our lab has explored a new approach we refer to as *subjective duration assessment* for gauging users' difficulties with web tasks, interfaces, and situations (Czerwinski, Horvitz & Cutrell, 2001). The approach, adapted from a finding described in the interruption literature in psychology, centers on the use of time estimation to characterize performance. We introduce a metric, named *relative* *subjective duration* (RSD) that provides a means for probing the difficulty that users have with performing web tasks—without directly asking users about the difficulty and invoking any demand characteristics that explicit probing might engender. Although there are several applications of RSD, we focused on the use of this measure to probe users' experiences with new website designs without directly asking them for feedback. Direct assessment of satisfaction has been found to be frequently confounded by an inherent bias toward the positive end of the scale. RSD is based on a discovery nearly seventy-five years ago showing that when engaging tasks are interrupted, participants tend to overestimate how long those tasks take when compared to actual task times. Conversely, tasks that are completed tend to be underestimated in terms of the overall task times. We explored the value of time estimation as a metric for evaluating task performance in web design. Our hypothesis was that participants would overestimate the duration of activity on tasks that are halted before completion because users were not able to complete them on their own. In contrast, we believed users would underestimate the duration of activity on tasks completed successfully. A user study of a well-known Internet browser explored the efficacy of the time estimation metric against both satisfaction and actual performance metrics during common web tasks. Our results showed that using subjective duration assessment can be a valuable tool for web user interface research. It circumvents the need to explicitly ask users for satisfaction ratings, and therefore is a better, implicit measure of web satisfaction that is more positively correlated with actual performance metrics. We hope to see more novel web interaction metrics emerge from laboratories doing advanced technology development over the years to come.

BEYOND THE DESKTOP—FUTURE WEB INTERFACES

We conclude this chapter with a selective glimpse at advanced web technology from our labs. The projects described below are currently being developed for deployment on the web over the next 2 to 5 years and are targeted towards making the web more useful for the mobile web user. Note that this technology is still so new and evolving so rapidly that it is yet unclear whether or how basic cognitive principles should or could be applied. Our approach has been to use a combination of iterative test and design in parallel with basic research on human interaction in order to shed more light on these matters. First we will discuss research on an audio feedback mechanism that will be useful for today's websites but will become more important for supporting mobile users. Then we will discuss current projects for supporting accessing the web over telephony and multimodal devices.

Non-Speech Audio

Ease of navigation is one of the most important elements of a good website, and the only tool we currently use to support navigation is visual feedback. We attempted to augment visual feedback with non-speech audio feedback on one of our websites. In previous work with this site we learned that people would tend to get stuck between the second and third levels of this website when looking for information and would forget to research the home page. We wanted to use audio in an unobtrusive way to let people know how deep they were in the website. Users didn't need to do anything extra to hear the signal; it would play whenever a link was selected. There were two dimensions coded in the signal. The first dimension of the signal provided depth information by pulsing the number of levels of depth of the selected item; the top level received one pulse, the second level item received two pulses, the third level item received three pulses. The second dimension coded into the audio was direction of movement; selecting an item on a lower level elicited a decreasing pitch while selecting an item on a higher level elicited an increasing pitch. It was expected that users would use this information to help navigate around the hierarchy and to not get trapped in an incorrect sub-section of the site. Unfortunately, there was no indication that the auditory feedback improved performance in the search task, the reaction times for finding particular items in the website were not statistically different between the nonaugmented site and when the audio feedback was added to the site. Similarly, the users did not reliably rate the website more appealing when the audio feedback was present. In future studies we plan to examine different kinds of non-speech audio to try to make the auditory feedback more understandable, and to also examine increased amounts of experience with the audio feedback. While on desktop devices non-speech audio will be redundant with visual feedback, it will become a crucial form of feedback when we try to provide make the web accessible over a telephone.

Telephony

WIMP (Windows, Icons, Menus and Pointers) interfaces on desktop computers have been very successful and are the most frequent way people currently access the web. In the future there will be a greater variety of devices used to connect to the web and the desktop input techniques will not scale well. One such device is the common phone. The phone is by far the world's most ubiquitous communication device and by opening up the web to anyone with a phone greatly increases the number of people who will have access to the web. Most phones do not support WIMP interfaces, content must be accessed with either key presses or speech recognition, and content must be displayed with either recorded speech for static content or text-to-speech (TTS). There are already a couple of commercial services for making web content available over the phone and others that are currently under development. Recently we compared three telephony interfaces that allow users to access stock reports, weather reports, and personal information content like email. These interfaces worked by presenting the user a list of the services that the system provided and the user would select the service by saying its name. Once in the system, the appropriate system features would be described and the user would again chose between options until the user drilled down far enough to reach the target information. Users were very successful with all three interfaces at accessing information like the local weather or the stock price for a particular company, but email systems are not as successful. Users had no problem navigating through their email box with commands like next and read, but these commands were not sufficient for navigating though more than a small handful of messages. There are several challenges for accessing the web over the phone. The current state of text-to-speech (TTS) technology is that it is good enough for people to clearly recognize all the words said, but that it does not sound natural enough for people to listen to for long periods of time. Second, telephony interfaces need to allow users the flexibility to change their mind at any time. Users expect to be able to stop the system at any time to do something different, and the systems currently are not always good at allowing barge-in. Barge-in will have to be universally available and feedback on the new action must be quick. Third, telephony systems need to figure out how to present data that are more complex. This is one area that non-speech audio can help, but the content itself will have to support the user's tasks. Reading through five emails one at a time is possible, but users do not want to read fifty emails like that. Users will not read email over the phone in the same way that they do on WIMP interfaces and we need to consider how their usage will differ. Task analysis will take on even greater importance for these devices. Next we turn to an exciting new kind of interface that is coming a little after telephone portals. These interfaces will present information on a display, and allow the user to use speech for navigation and for text entry.

Speech & Mobile

For the web to become truly ubiquitous, it must be accessible at anytime from anyplace. The current generation of PDAs (Palm and PocketPC devices) are the first computing devices that are truly portable enough to bring anywhere. Unfortunately these devices provide an unsatisfactory web experience because navigation and text entry are much too difficult to access on a small screen and without a keyboard. MiPad, a prototype used in our labs, uses a combination of speech recognition and natural language capabilities to allow users to be more productive with PDAs (Hwang et. al., 2000). MiPad lets users issue commands like "Schedule a 30 minute meeting with

Bill Gates for Friday at 2 in the afternoon", which will cause a new appointment to be opened and all the appropriate fields filled in. Users can also write new email messages with MiPad by dictating the text they want sent. User studies show that users can schedule appointments and transcribe emails with the MiPad speech interface reliably faster than with a soft-keyboard on the same PDA. A key challenge here is correcting speech recognition errors. Speech recognition accuracy is in the neighborhood of 95%, which means that 1 word out of every 20 needs to be corrected. It is this correction time that keeps speech recognition from being more efficient than typing on a desktop machine. We need to develop strategies for reducing the time for correction. In a recent study in our labs we found that using speech-only to correct speech errors is terribly inefficient because they lead to many additional errors. Using either an alternates list to correct errors or a multimodal correction technique with pen to select and speech to redictate is better than speech-only.

Recently we have turned our attention from developing applications that use speech on a PDA to developing the framework that will enable anyone to turn their regular websites into speech enabled websites. Developers will be able to write Speech Application Language Tags (SALT) (www.saltforum.org) to let people interact with websites with speech. SALT can work on any device from a simple phone to a desktop, but its sweet spot is certainly enabling mobile interaction. Being able to ask a website for flight times for Seattle to New York and seeing a list of flights is better than either using a phone and hearing a list of flights or having to use inefficient input methods on the PDA. In addition to using SALT for form-filling on the web, if it is combined with natural language technologies, users will be able to ask information and expect the website to handle the navigation rather than having to perform traditional navigation on a small screen. As we move forward with this we will face the challenges of efficiently correcting the speech recognition errors and developing the natural language interface that will be used to navigate both between and within websites.

CONCLUSION

We opened this paper with the premise that the design and evaluation of web sites is challenging. We further argue that basic results and theory can provide reasonable starting places for design, although perhaps experience in an application domain is as good. Theories and findings from psychology are not, however, as useful as one might hope. Getting the design right from the beginning is very hard even with these starting points, so any real world design will involve an ongoing cycle of design, evaluation, failure analysis, and redesign.

How do we move from the current state of affairs to a better symbiosis between basic research in cognitive science and web design in order to eradicate over-reliance on web designs that are "point" designs, or the examination of web features in complete isolation for simplistic tasks? We would like to propose moving basic research in cognitive science closer to the real needs of web designers and human-computer practitioners by studying more complex, real world problem domains. Examples ripe for investigation include such popular web domains as purchasing items on the web, or searching for information. We do, however, feel as though progress is being made. For instance, starting with the variables discussed in this chapter (short-term memory, visual attention dimensions, notifications and task switching, information retrieval, input and speech), the web practitioner can begin to see basic guidelines for design emerging from carefully controlled studies of the intended web behavior. It is our vision that increasingly more academic and industrial research partners will work toward common goals in building a better foundation for web design.

REFERENCES

Callaghan, T. C. (1990). Interference and Dominance in Texture Segregation. In *Visual Search*, Brogan, D., Ed., 81-87. Taylor & Francis, New York, New York.

Chen & Czerwinski, 1999;

Chen, H. & Dumais, S.T. (2000). Bringing order to the web: Automatically categorizing search results. In *Proceedings of Association for Computing Machinery's CHI 2000: Human Factors in Computing Systems*, Addison-Wesley, 145-152.

Cutrell, E., Czerwinski, M. & Horvitz, E. (2001). Notification, disruption and memory: Effects of messaging interruptions on memory and performance. In Hirose, M. (Ed.) *Proceedings of Human-Computer Interaction -- Interact 2001*, Tokyo, IOS Press, copyright IFIP, 2001, 263-269.

Czerwinski, M., Cutrell, E. & Horvitz, E. (2000). Instant messaging: effects of relevance and time. In S. Turner, P. Turner (Eds), *People and Computers XIV: Proceedings of HCI 2000, Vol. 2*, British Computer Society, 71-76.

Czerwinski, M., Horvitz, E. & Cutrell, E. (2001). Subjective duration assessment: An implicit probe for software usability. In Proceedings of Benyon, D. & Palanque, P. (Eds *People and Computers XV: Proceedings of HCI 2001, Vol. 2*, British Computer Society, *-*.

Czerwinski, M., van Dantzich, M., Robertson, G.G. & Hoffman, H. (1999). The contribution of thumbnail image, mouse-over text and spatial location memory to web page retrieval in 3D. In Sasse, A. & Johnson, C. (Eds), *Human-Computer Interaction--Proceedings of Interact '99, Edinburgh*, Scotland, IOS press, pp. 163-170.

Dumais, S.T., Cutrell, E. & Chen, H. (2001). Optimizing search by showing results in context. In *Proceedings of Association for Computing Machinery's CHI 2001: Human Factors in Computing Systems*, Addison-Wesley, 277-283.

Dumais, S.T. and Czerwinski, M. P. (2001). Building bridges from theory to practice. In *HCI International 2001, 9th Conference on Human-Computer Interaction*. Session on Contributions of Basic Research in Psychology to Human Factors and Human-Computer Interaction. August 2001.

Hudson, W. (2001). Designing for the grand illusion. In Association for Computing Machinery's *SIGCHI Bulletin*, November/December, 2001, 8 (http://www.acm.org/sigchi/bulletin/2001.6).

Hwang, X., Acero, A., Chelba, C.,Deng, L., Duchene, D., Goodman, J., Hon, H., Jacoby, D., Jiang, L., Loynd, R., Mahajan, M., Mau, P., Meredith, S., Mughal, S., Neto, S., Plumpe, M., Stery, K., Venolia, G., Wang, K., Wang, Y. (2000). MIPAD: A Multimodal Interactive Prototype, in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing.* Istanbul, Turkey, June, 2000.

Larson, K. & Czerwinski, M. (1998). Web page design: Implications of memory, structure and scent for information retrieval. In *Proceedings of Association for Computing Machinery's CHI '98, Human Factors in Computing Systems*, ACM press, 25-32.

Lim, R.W., & Wogalter, M.S. (2000). The position of static and on-off banners in WWW displays on subsequent recognition. In Proceedings of the IEA 2000/HFES 2000 Congress. San Diego, CA: Human Factors and Ergonomics Society, pp. 420-423.

Lohse, G.L. & Spiller, P. (1998). Quantifying the effect of user interface design features on cyberstore traffic and sales. In *Proceedings of Association for Computing Machinery's CHI '98: Human Factors in Computing Systems*, Addison-Wesley, 211-218.

Lynch, K. (1960). The Image of the City. Cambridge, Massachusetts: The MIT Press.

Passini, R. (1984). Wayfinding in Architecture. New York: Van Nostrand Reinhold.

Robertson, G. , Czerwinski, M., Larson, K., Robbins, D., Thiel, D. & van Dantzich, M. (1998). Data Mountain: Using Spatial Memory for Document Management. In *Proceedings of UIST '98, 11th Annual Symposium on User Interface Software and Technology*, pp. 153-162.

Tavanti, M. & Lind, M. (2001). 2D vs. 3D, Implications on spatial memory. In Proceedings of *IEEE's Information Visualization*, 139-145.

Tiernan, S., Cutrell, E. Czerwinski, M. & Hoffman, H. (2001). Effective notification systems depend on user trust. In Hirose, M. (Ed.) *Proceedings of Human-Computer Interaction -- Interact 2001*, Tokyo, IOS Press, copyright IFIP, 2001, 684-685.

Triesman, A. (1985). Preattentive Processing in Vision. *Computer Vision, Graphics, and Image Processing 31*, 156-177.

Tullis, T.S. (1998). A method for evaluating web page design concepts. In *Proceedings of Association for Computing Machinery's CHI '98: Human Factors in Computing Systems*, Addison-Wesley, 323-324.

Vinson, N.G. (1999). Design guidelines for landmarks to support navigation in virtual environments. In *Proceedings of Association for Computing Machinery's CHI '99: Human Factors in Computing Systems*, Addison-Wesley, 278-285.


Corbis - The Place for Pictures Online - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://www.corbis.com/
Corbis
The Place for Pictures Online™
Welcome to the Web's largest collection of high-quality images for professional and personal use—featuring the world's top photographers and most beloved artists. Just click a square to follow your interest.
Shopping
Professional Licensing
BizPresenter™
About Corbis
Stunning imagery with a Euro flair. Check out our latest R-F catalog.
Limited editions from the Bettmann Archive: See The Living Lens.
Find great images for PowerPoint(r) or Web use at BizPresenter.
Balloons Floating Over U.S. Capitol Dome © Joseph Sohm/JS1000596
© 2001 Corbis Corporation. All rights reserved.
Done
Internet

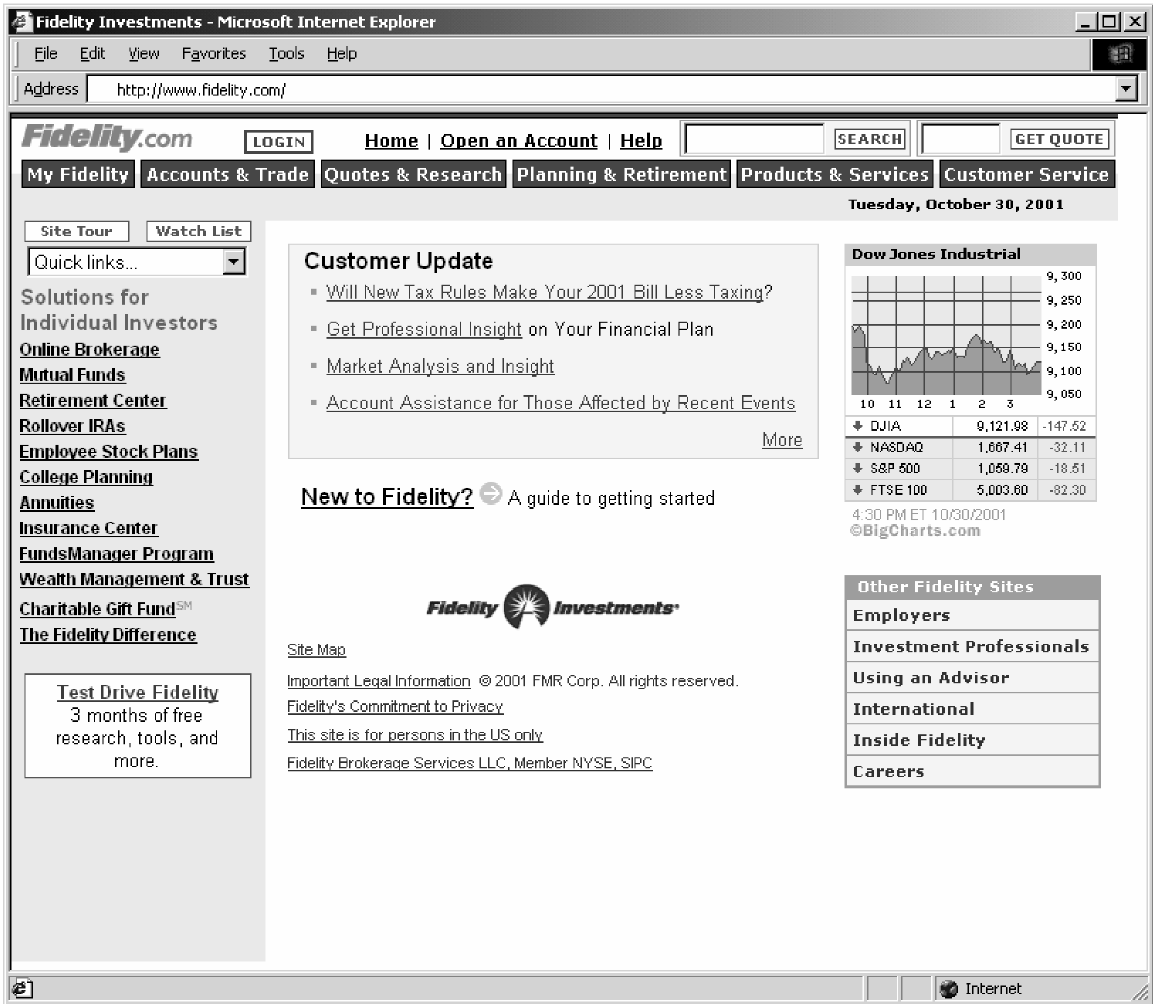
Fidelity Investments - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://www.fidelity.com/
Fidelity.com
LOGIN
Home | Open an Account | Help
SEARCH
GET QUOTE
My Fidelity
Accounts & Trade
Quotes & Research
Planning & Retirement
Products & Services
Customer Service
Tuesday, October 30, 2001
Site Tour
Watch List
Quick links...
Solutions for Individual Investors
Online Brokerage
Mutual Funds
Retirement Center
Rollover IRAs
Employee Stock Plans
College Planning
Annuities
Insurance Center
FundsManager Program
Wealth Management & Trust
Charitable Gift Fund
The Fidelity Difference
Test Drive Fidelity
3 months of free research, tools, and more.
Customer Update
Will New Tax Rules Make Your 2001 Bill Less Taxing?
Get Professional Insight on Your Financial Plan
Market Analysis and Insight
Account Assistance for Those Affected by Recent Events
More
New to Fidelity? A guide to getting started
Fidelity Investments
Site Map
Important Legal Information © 2001 FMR Corp. All rights reserved.
Fidelity's Commitment to Privacy
This site is for persons in the US only
Fidelity Brokerage Services LLC, Member NYSE, SIPC
Dow Jones Industrial
9,300
9,250
9,200
9,150
9,100
9,050
10 11 12 1 2 3
DJIA 9,121.98 -147.52
NASDAQ 1,667.41 -32.11
S&P 500 1,059.79 -18.51
FTSE 100 5,003.60 -82.30
4:30 PM ET 10/30/2001
©BigCharts.com
Other Fidelity Sites
Employers
Investment Professionals
Using an Advisor
International
Inside Fidelity
Careers
Internet

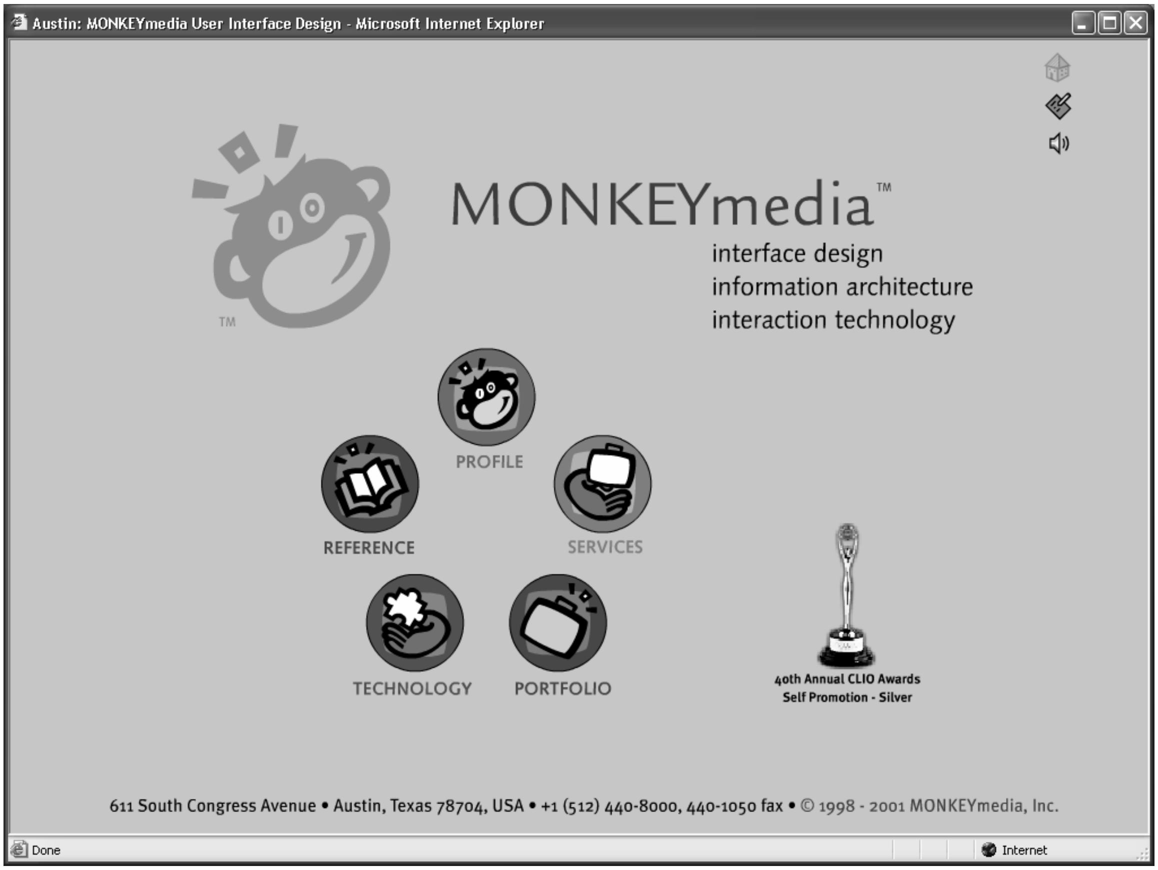

Insert SWISH EXAMPLE Here