HUMAN FACTORS AND WEB DEVELOPMENT

Edited by
Chris Forsythe
Eric Grose
Julie Ratner

WESTERN MICHIGAN UNIVERSITY
3 1141 01289 7752

DATE DUE

AUG 0 7 2000
SEP 0 1 00
OCT 1 8 '00
OCT 26 '00

ALL ITEMS ARE SUBJECT TO RECALL
AFTER 2 WEEKS

Human Factors and Web Development

TK
5105.9
.H86
1998
(ERC)

Human Factors and Web Development

Edited by

Chris Forsythe
Sandia National Laboratories

Eric Grose
Sandia National Laboratories

Julie Ratner
The MathWorks, Inc

LEA
LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
1998 Mahwah, New Jersey London

The final camera copy for this work was prepared by the editors, and therefore the publisher takes no responsibility for consistency or correctness of typographical style. However, this arrangement helps to make publication of this kind of scholarship possible.

Copyright © 1998 by Lawrence Erlbaum Associates, Inc.
All rights reserved. No part of this book may be reproduced in any form by photostatic, microfilm, retrieval system, or any other means, without the prior written permission of the publisher.

Lawrence Erlbaum Associates, Inc., Publishers
10 Industrial Avenue
Mahwah, NJ 07430

Library of Congress Cataloging-in-Publication Data

Human Factors and Web Development / Chris Forsythe, Eric Grose & Julie Ratner.
p. cm.
Includes bibliographical references and index.
ISBN 0-8058-2823-0 (c : alk. Paper). -- ISBN 0-8058-2824-9 (p : alk. Paper)
1. Web sites—Design. 2. Web sites—Psychological aspects. 3. World Wide Web (Information retrieval system) – Psychological aspects. 4. Human computer interaction. I. Forsythe, Chris. II. Grose, Eric. III. Ratner, Julie.
TK5105.888.H86 1997
025.04—DC21 97-15656
CIP

Books published by Lawrence Erlbaum Associates are printed on acid-free paper, and their bindings are chosen for strength and durability.

Printed in the United States of America
10 9 8 7 6 5 4 3 2 1

Contents

About the Contributors vii

Preface ix

1 Introduction 1
Deborah J. Mayhew

Part I: PERSPECTIVES FROM PSYCHOLOGY

2 The Use of Investigatory Responses as a Measure of Learning and Memory 17
Charles P. Seltzer

3 Visual Information Processing on the World Wide Web 25
William Marks & Cynthia L. Dulaney

4 Discourse Process and Its Relevance to the Web 45
Joseph P. Magliano, Michelle C. Schleich, & Keith K. Millis

5 Human Navigation 63
Leslie A. Whitaker

Part II: WEB USER POPULATIONS

6 Children's Online Environments 75
Allison Druin & Mark Platt

7 Designing Web Pages and Applications for People With Disabilities 87
Lila Laux

8 The World Wide Web as a Teaching Resource 97
Peter Burden & Jenny Davies

9 Easing the Learning Curve for Novice Web Users 111
Julie Ratner

Part III: WEB DESIGN GUIDELINES AND DEVELOPMENT PROCESSES

10 Using Web and Traditional Style Guides to Design Web Interfaces 121
Eric Grose, Chris Forsythe, & Julie Ratner

11 Page Design Guidelines Developed Through Usability Testing 137
José A. Borges, Israel Morales, & Néstor J. Rodríguez

12 Human Factors Methodology for Designing Web Sites 153
Pawan Vora

Part IV: WEB RESEARCH AND DEVELOPMENT

13 Web User Interface Development at Oracle Corporation 175
Anna M. Wichansky & George Hackman, Jr.

14 Web Usability Research at Microsoft Corporation 189
Amy Kanerva, Kevin Keeker, Kirsten Risden, Eric Schuh, & Mary Czerwinski

15 Creating Content for Both Paper and the Web 199
Richard C. Omanson, Gavin S. Lew, & Robert M. Schumacher

16 The Ten Golden Rules for Providing Video Over the Web or 0% of 2.4M (at 270k/sec, 340 sec remaining) 207
Chris Johnson

Part V: COLLABORATION AND VISUALIZATION

17 Graphics Design on the Web 225
Eric N. Wiebe & Julie E. Howe

18 Collaborative Interfaces for the Web 241
Saul Greenberg

19 A Zooming Web Browser 255
Benjamin B. Bederson, James D. Hollan, Jason Stewart, David Rogers, David Vick, Laura Ring, Eric Grose, & Chris Forsythe

Author Index 267

Subject Index 275

About the Contributors

Benjamin B. Bederson University of New Mexico

José A. Borges University of Puerto Rico - Mayagüez Campus

Peter Burden University of Wolverhampton

Mary Czerwinski Microsoft Corporation

Jenny Davies University of Wolverhampton

Allison Druin University of New Mexico

Cynthia L. Dulaney Xavier University

Chris Forsythe Sandia National Laboratories

Saul Greenberg University of Calgary

Eric Grose Sandia National Laboratories

George Hackman, Jr. Oracle Corporation

James D. Hollan University of New Mexico

Julie E. Howe North Carolina State University

Chris Johnson University of Glasgow

Amy Kanerva Microsoft Corporation

Kevin Keeker Microsoft Corporation

Lila Laux US WEST Communications

Gavin S. Lew Ameritech

Joseph P. Magliano Northern Illinois University

William Marks The University of Memphis

Deborah J. Mayhew Deborah J. Mayhew & Associates

Keith K. Millis Northern Illinois University

Israel Morales University of Puerto Rico - Mayagüez Campus

Richard C. Omanson Ameritech

Mark Platt University of New Mexico

Julie Ratner The Mathworks, Inc

Laura Ring Sandia National Laboratories

Kirsten Risden Microsoft Corporation

Néstor J. Rodríguez University of Puerto Rico - Mayagüez Campus

David Rogers University of New Mexico

Michelle C. Schleich Northern Illinois University

Eric Schuh Microsoft Corporation

Robert M. Schumacher Ameritech

Charles P. Seltzer Essex Community College

Jason Stewart University of New Mexico

David Vick University of New Mexico

Pawan Vora U S WEST Communications

Leslie A. Whitaker University of Dayton

Anna M. Wichansky Oracle Corporation

Eric N. Wiebe North Carolina State University

Preface

Publication and mass distribution of information products (e.g., books, magazines, newspapers) has traditionally been limited to the relatively few private and public entities with the necessary resources and established distribution networks. This situation has led to the development of various specialists with the knowledge and skills to deliver quality products. Across the population in general and in large corporate and government organizations, Web-based technologies have expanded the number and range of people capable of distributing information to international audiences, until information can now reach almost anyone with access to a personal computer.

Over the past 2 decades, standards for the usability of human-computer user interfaces (UIs) have emerged. To assure the quality of human-computer interfaces, practitioners trained in Human Factors have become routine participants in computer software and hardware development processes. For most Web-based materials, the developers are neither trained in Human Factors, nor have access to Human Factors expertise. The Human Factors profession, with its unique knowledge and practices, has only had a marginal impact on the World Wide Web (WWW) and Web developers.

Traditional Human Factors offers many insights about achieving usable, Web-based interfaces. As Web-based products have begun to permeate the lives of more and more computer users, Human Factors practitioners have begun to contemplate the human-computer interface issues arising with the WWW and have begun to experiment with solutions for improving the usability of Web products. The growth of the WWW, however, has occurred at such a rapid pace that, to date, virtually no comprehensive sources have emerged as guides for the design of usable Web products.

This book includes the perspectives of a broad range of experts in Human Factors, Cognitive Psychology, and Web development and reports research addressing issues associated with the design and usability of Web products. In doing so, the book simultaneously seeks to fill two voids: First, it makes the expertise, knowledge, and practices of Human Factors available to the Web development community. Second, it provides a comprehensive source for Human Factors practitioners responsible for the design of usable Web-based products.

Chris Forsythe
Sandia National laboratories

1 Introduction

Deborah J. Mayhew
Deborah J. Mayhew & Associates

HUMAN FACTORS AND THE WEB

The World Wide Web (WWW) is still relatively new. It began in 1989, the brainchild of Tim Bernes-Lee at the European Organization for Nuclear Research (CERN), to address the need for a collaborative tool to help scientists share knowledge. It was then offered to the entire Internet community to experiment with, and thus the WWW was born (Horton, Taylor, Ignacio, & Hoft, 1996, p. 13).

What is the WWW? It is a huge, ever-growing collection of hyperlinked documents created by independent authors, stored on computers known as Web servers, and made accessible to anyone over the Internet via software applications called browsers and various search engines accessible through browsers. Browsers provide a relatively user-friendly interface (at least compared to trying to use the Internet in the days before browsers) for navigating and viewing the total information space represented by all sites on the WWW. Thus, whereas the Internet is an electronic communications network that generally supports various kinds of communications including E-mail, the WWW represents a repository of public information created by the public and accessible to the public via the Internet.

By analogy, the Internet is an electronic pony express, providing an electronic means to move information from one physical location to another; E-mail is an on-line post office/postal service, providing an on-line process for using the pony express (the Internet) to send and receive private mail; the WWW is an on-line library, storing and providing public access to a huge body of information; and search engines and browsers are partly an on-line Dewey Decimal classification system/card catalog, providing an on-line method for finding and viewing information stored on the WWW over the Internet.

Traditional public libraries, however, are very different from the WWW. On the WWW, anyone can publish a book (a Web site) and add it to the library (the WWW). No publishing companies screen and edit books for publication, and no librarians decide what and what not to add to a library's collection. Because screening, editing, publishing and purchasing intermediaries do not exist in the world of the WWW, prospective authors themselves must judge what content is useful and will be used, and what presentation is effective. WWW users must

wade through the overwhelming morass of unscreened, unedited published information to find useful, credible, well-presented information. One of the WWW's great powers is its openness. It represents the ultimate in freedom of speech and makes huge, geographically dispersed audiences easily available to any aspiring author or publisher at very low cost. One of its great weaknesses, on the other hand, is its lack of quality control, both for content and for presentation. This role has traditionally been assumed by publishing companies and libraries.

If it somehow became possible and easy for anybody to self-publish a book and distribute it to any library, two new services would probably emerge – one to help author/publishers produce high-quality publications and another to help library users find, screen, and evaluate publications. In the case of the WWW, which makes self-publication easy, search engines and browsers help users find publications, and a well-established and completely applicable service is available to help prospective author/publishers produce high-quality publications - the field of software usability engineering.

Software usability engineering, or software human factors, is a relatively new field, less than 20 years old. During this time, it has produced useful methods and guidelines for designing usable software—methods and guidelines that are directly applicable to the design of Web sites.

There are many reasons why an individual or organization might decide to publish a Web site. These include (but are certainly not limited to) attracting business or offering a service; sharing a special skill, knowledge or perspective; giving specific answers whenever users need them; shrinking distances and reducing isolation of all kinds; supporting learning and intellectual stimulation and pursuing personal interests (Horton et al., 1996, pp. 3–6)

Because these reasons are very broad, the range of Web site author/publishers is also broad. Many are unaware of the field of usability engineering and what it can offer them in their efforts to produce a useful, usable Web site. One purpose of this book is to introduce potential or current Web author/publishers to this field to show its relevance to Web design.

Reasons for browsing the WWW are just as broad and varied, including the obvious, usual reasons such as finding information about products and services, researching a topic, accessing up-to-date information about such things as stock markets and finding people or organizations with similar interests. Some less usual reasons include finding the birth parents of an adopted child, getting information on wanted criminals, getting on-line access to comics, and finding information about colleges and universities.

Many if not most WWW users are discretionary; many have had relatively little computer experience, and some may be computer-phobic; many are only occasional users; and many may not speak English as their native language (the Web is indeed worldwide). Thus, usability issues in Web design seem even more critical than they are for traditional software application design.

Consider the following interesting statistics which have undoubtedly changed dramatically since this writing, but can serve as examples:

- 7.7% of U.S. adults use the Internet (http://www.openmarket.com/intindex/96-07.htm).
- The average Internet user's age is 35 years: the average household income of Internet users is $69,000 (Foster, 1996).
- 70% of Internet users are male, 76% are U.S. residents, 52% are Microsoft Windows users, 26% are Apple Macintosh users, and 77% browse the WWW at least once a day (Foster, 1996).
- There are over 30 million Internet users (Horton et al., 1996, p. 9).
- The Internet is growing at 20% per month (Foster, 1996).
- The WWW was born in 1989—only 8 years ago (Horton et al., 1996).
- 82% of time spent on the Internet is spent Web browsing (Foster, 1996).
- 18 million people used the WWW in a three month period in 1995 (http://www.openmarket.com/intindex/96-01.htm).
- There are almost a million Web sites (Nielsen, 1996).
- The estimated volume of sales generated by the World Wide Web in 1998 is $46 billion (http://www.openmarket.com/intindex/96-07.htm).
- Of the more than 10 million U.S. businesses listed in the Central Source Yellow Pages, 15,379 are registered as having Web sites in the Commercial Sites Index, and an average of 73 businesses are added each day (http://www.openmarket.com/intindex/96-01.htm).
- 66% of television advertisements refer directly or indirectly to Web sites (Foster, 1996).
- Readership of Jakob Neilsen's monthly Web column on Web UI Design (www.sun.com/columns/alertbox.html) has been growing 5,000% per year (Nielsen, 1996).

Clearly the Web is here, and here to stay.

Initially and until very recently, Web sites have been read-only information sites, and the only interaction supported on Web sites has been controls (links) for navigating to other parts of a site or to other sites. In addition, tools for developing Web sites have been limited to providing ways to define content; presentation of content has been mostly outside the control of individual Web site authors, and under the control partly of browser applications and partly of users through the browser. Tools are now emerging to support the development of interactive Web sites that resemble traditional types of interactive software and to provide support for design and presentation to Web authors.

Even when designing read-only Web sites, designers can draw on usability issues and principles from the field of traditional software usability. Much traditional software user interface (UI) design involves effective presentation of information and the design of effective navigational aids and methods. These issues are common to both traditional software application and traditional Web site design. As Web sites become more interactive and new tools and technologies provide more design options to developers, the principles, guidelines, and methods of software usability engineering become even more directly applicable.

A random tour around the WWW, however, shows that the knowledge and wisdom collected in the field of software usability engineering over the last 20 years is generally not being applied to Web site design.

Some useful distinctions between Web design and traditional software design include the following (Nielsen, 1996):

- On the Web, response times are generally slower and less predictable than on traditional software, and this places significant limits on practical file sizes and interaction in general.
- The WWW has fewer restrictions on style, and no interaction guidelines have yet emerged comparable to the interaction style guides available to traditional software developers.
- On the WWW, designer-developers have much less control of the appearance of information; they control primarily only content and overall structure.
- Web users, unlike most users of traditional software, are mainly discretionary users, a fact that increases the importance that Web software be "walk-up-and-use."
- The WWW is huge, much larger than any application, with fluid boundaries between sites. It has an increased need for navigational support and sense of place in the UI.

These differences are not qualitative but are more a matter of degree. They hearken back to the days of mainframe computers and dumb terminals, when designers were similarly constrained by poor performance, the absence of style guides and major restrictions on display capabilities. Today many applications are intended for an audience of discretionary, generally computer-naive users. Similarly, automation today is moving away from discrete, unintegrated applications, toward suites of separate but tightly integrated applications that share databases and common UIs. As Web development tools become more capable, traditional UI issues arise, such as:

- The importance of good conceptual model design to communicate structure, support user models, and facilitate effective navigation (Nielsen, 1996).
- The importance of context information to maintain a sense of place, both global (i.e., on the WWW) and local, (i.e., in a site) (Nielsen, 1996).
- The importance of good screen design.
- The importance of effective feedback.
- The importance of user-centered design and usability testing (Web design is interaction design, just as other types of software design) (Nielsen, 1996).

Thus, Web design does not call for a new field to support designers. We need only look to the field of software usability engineering to see the strong analogies between traditional and Web software. This book reinforces this view by offering examples of the application of usability engineering methods and principles to Web design.

ABOUT THIS BOOK

The first part of the book is titled Perspectives From Psychology; the authors have drawn on their expertise in various subfields of psychology to sketch implications for Web design.

In chapter 2, "The Use of Investigatory Responses as a Measure of Learning and Memory," Seltzer reviewed and summarized research on the role of exploratory, non task-driven behavior in learning and drew implications for Web use. Research indicates that many animals seem instinctively to engage in environmental investigation without any apparent specific motivation and that their familiarity with a general terrain assists them in solving problems that arise later, such as escaping from predators. Seltzer suggested an interesting analogy to users on the WWW. Users who surf the Web without a particular goal are learning information, both about the WWW as an environment and real content, which may become useful to them in later contexts. An implication from this analogy from the viewpoint of UI design and the WWW, is that exploratory, non-goal-oriented navigation on the WWW should be supported and facilitated by Web design. Here, perhaps, is an area where an interface design issue on the WWW is truly different from design issues typically encountered when designing traditional applications, which usually are meant to support goal-oriented tasks. Perhaps there is a connection between UI design for games and for the WWW.

In chapter 3, "Visual Information Processing," Marks provided a discussion of human visual processing, with its constraints and capabilities, as it applies to Web use and thus design. Marks described aspects of visual processing including pattern recognition, motion perception, reading, visual searching, and visual attention, and drew implications for designing effective and usable Web displays.

Chapter 4, "Discourse Process and Its Relevance to the Web," by Magliano, Schleich and Millis, includes a review of the literature on discourse processing, which is the study of how people produce and understand language. Readers gain insight into how they process, comprehend, and remember different kinds of prose, and the authors have offered four principles for designing prose to communicate most effectively, ideas that can be applied to Web design.

In chapter 5, "Human Navigation," Whitaker addressed the problem of designing for easy and effective navigation in and across Web sites by reviewing research into human navigation in other contexts. She distinguished between navigation for two different purposes—to achieve a specific goal, and simply to explore with no specific goal in mind—and has noted that people use different

strategies depending on their purposes. She also distinguished between navigation through structured (manmade) and natural environments and has pointed to different user strategies in these different contexts, to apply the findings to Web design. She studied human navigation in such contexts as navigation by aircraft pilots, navigation by people in unfamiliar and familiar cities, and navigation in off-road, natural environments. As a summary Whitaker has listed 10 concrete design principles for navigation on the Web.

Part II, "User Populations", offers a look at design issues relevant to specific Web user populations. In chapter 6, "Children's On-line Environments," Druin and Platt focused on the appeal and utility of the Internet to a unique user group - children. In several experiments, children used the Internet for educational purposes. In one case, children at an after-school program in a rural community were introduced to the Internet to communicate with health care workers about issues in health, nutrition, and medicine. In another case, children at an after-school program in a low-income area where support for music and art had recently been cut, were introduced to the Internet as part of a program that supported exploring music. Although Druin and Platt cited exciting examples of the benefits for children available through the Internet, from these cases they also drew several conclusions about future developments that could better support the needs of this important user group - support for drawing as a communication medium (kids would rather draw than type), tools that better support collaborative work (kids love to work together), and UIs that are easy for children to learn.

This last point struck me as particularly interesting, because of the example from which it was drawn. Druin and Platt described that they had a hard time communicating to children the difference between E-mail and the WWW. This distinction was not at all clear to children, and they found some creative ways of explaining it through exercises with balloons, exploiting the notion of a metaphor to support learning. It struck me that many adults also have trouble grasping the distinction among the Internet, E-mail, and the WWW. As Druin and Platt noted, there is a flaw in the design of these tools—there is no clearly and visually communicated conceptual model. The lack of a clear conceptual model is a general flaw that can arise in the design of any software, not just Web software. There are information-gathering and design approaches available to design good conceptual models for software. Again, this general UI issue also arises in the context of Web software.

Chapter 7, "Designing Web Pages and Web Applications for People With Disabilities," by Laux, describes another user group for whom the Internet holds great potential benefits—people who are ill, disabled and older. Laux offered simple, concrete, and practical advice on how to design for users who are visually, audibly or motor challenged, and on how to test Web interfaces to make sure that they meet the usability needs of this potentially large class of users.

In chapter 8, "The World Wide Web as a Teaching Resource," Burden and Davies have discussed the current and future potential use of the WWW by students in an educational setting who use the WWW as an alternative or additional source of information beyond a traditional library. Such issues as the

difficulty of finding appropriate resources and the quality, permanence, and ease of access of these resources are addressed.

Chapter 9, "Easing the Learning Curve for Novice Web Users," by Ratner, also looks at the use of the WWW in an educational setting, but from the perspective of using the WWW to deliver curriculum, rather than as an alternative to using a library. From her research on novice Web users in this context, she, like other authors in this volume, reported issues in using browser functionality which seem to point to the lack of a good conceptual model in the design of browser UIs. This problem is also a classic one in UI design, and Ratner has noted that novices have particular problems when trying to become familiar with the WWW. Ratner has offered advice on how to provide effective orientation and training to novice Web users in an educational setting, in view of the current design of Web browsers. The design recommendations of other authors in this volume have also suggested ways to make the WWW more usable to novices through redesign rather than training.

Part III offers chapters that focus on different aspects of Web design. In chapter 10, "Using Web and Traditional Style Guides to Design Web Interfaces," Grose, Forsythe, and Ratner have reported on their investigation of available Web style guides, and have compared them to traditional software UI style guides. They found significant differences in the content of the two types of style guides, with only about 20% of the content of traditional style guides being covered in style guides tailored to the WWW. In particular, they noted that available Web style guides tend to focus on technical constraints and solutions to them, rather than on usability issues. In trying to understand the source of the difference, they pointed out that Web authors (and Web style guide authors) are a wide and varied group relative to traditional software developers and many of them simply have not been exposed to the field of usability engineering and the available traditional style guides.

After surveying several Web sites and analyzing the most common usability flaws, Grose et al. developed a short Web style guide (about 25 guidelines) that, if followed, would minimize the most common usability flaws in Web sites. They pointed out, however, that in Web sites there are still many violations of basic usability that are not addressed in their short Web style guide. In addition, they observed that, as in traditional software UI design, many design issues are dependent on knowing specifics about users and tasks and these issues cannot be addressed in a short and relatively simple style guide. They concluded that simple Web-specific style guides, although helpful, cannot substitute for traditional usability engineering expertise and for comprehensive style guides during Web development. Usability engineers must take traditional style guides, which have been rigorously developed using usability engineering techniques, and tailor them to be Web specific - not by shortening them or even by changing the basic principles very much, but by presenting new examples of the realization of these general principles on actual Web sites. Engineers probably need to make such style guides available directly over the WWW, because, as Grose et al.

pointed out, many if not most Web authors do not look beyond the WWW itself for guidance on how to develop Web sites.

In chapter 11, "Page Design Guidelines Developed Through Usability Testing", Borges, Morales, and Rodriguez observed that many traditional UI design flaws can be found on Web sites; they have offered a set of simple guidelines for shaping the design of Web sites and minimizing usability problems. The guidelines were developed on the basis of the usability problems most commonly found in a survey of Web sites, and tested by having designers use the guidelines to redesign some Web pages. When the performance of users using those pages was tested, the authors found significant improvements in usability on the pages redesigned according to the guidelines, compared to the original pages. Although the guidelines were tailored to the specific characteristics of Web site design, they are all instances of general principles that apply to the design of any software interface, such as good grouping, clear navigational pathways, clear terminology, and consistency.

Chapter 12, "Human Factors Methodology for Designing Web Sites," by Vora, describes the usability engineering process applied at US WEST Communications for the development of Web sites. This methodology is almost identical to usability engineering processes that have evolved over the past 20 years to support traditional software design and development and includes such traditional usability engineering tasks as user profiling, task analysis, usability goal setting, identifying platform constraints and capabilities, conceptual model design and screen standards design (style guide development), and iterative usability testing. Vora has also highlighted areas of this process in which Web design presents some unique challenges. These areas include: the need to address the requirements of Web authors and owners as well as Web users; designing for communication between people via Web sites; designing for the fact that users may enter a site through any page, not just a home page; designing for lower bandwidth and slower response times; designing for exploratory versus goal-oriented usage; designing for and testing on a wide, unpredictable variety of platforms and browsers; designing for ease of constant content updates without affecting usability; and designing to allow user feedback directly from a Web site. Vora has offered specific insights into these issues in the context of an overall methodology for Web design and development.

Whereas section I offered psychology research that can potentially be applied to Web design, Part IV provides chapters that described Web research exploring Web-specific design issues.

Chapter 13, "Web User Interface Development at Oracle Corporation," by Wichansky and Hackman, describes how Oracle's usability staff participated in the development of several of Oracle's current Web interfaces. The authors have discussed a corporate vision in which individual user workstations have minimal local storage and users access software applications running on Web servers. This process, they suggested, would make software more universally accessible and would reduce the cost of providing individual users with software applications. With this vision in mind, Oracle has begun "Web-enabling" some of their

standard software packages, and the usability staff has participated in the redesign of these packages to run on the WWW.

Wichansky and Hackman demonstrated, through their description of the usability work applied to the development of Oracle Web sites, that traditional software usability methods and principles are appropriate and applicable to Web development. They described performing user and task analyses, prototyping, and usability testing for Web interfaces under development, and they referred to classic design issues that arose in their testing, such as power versus ease of learning, meeting user expectations by exploiting current user models, good conceptual model design, the importance of context information and navigational support, and response time.

Chapter 14, "Web Usability Research at Microsoft Corporation," by Kanerva, Kecker, Risden, Shuh, and Czerwinski, described the usability work applied by the Interactive Media Group at Microsoft during development of Microsoft's Web products, which include not only their browser, but also over 50 content sites. This group has developed design principles and guidelines specific to Web development which can be used by Web design-development teams; conducted research to establish objective measures of Web site usability; designed and tested interface mechanisms to support search behaviors on the WWW, and conducted research to study users' mental models of information to support optimal organization of Web sites. In addition, they have studied what makes Web sites appealing, engaging, and motivating to users of different ages and genders, in this case going beyond traditional definitions of usability and exploring the entertainment and marketing potential of Web sites. The group has adapted and tailored traditional usability engineering techniques such as task analysis, usability goal setting, guidelines development, heuristic evaluation, and usability testing to suit the specific needs of Microsoft's Web development organization. The application of these techniques is described in some detail in this chapter.

Chapter 15, "Creating Content for Both Paper and the Web," by Omanson, Lew, and Schumacher, reported studies on how organizations approach the problem of Web design. The authors described several strategies and their relative success, from putting existing paper-based materials online, to designing explicitly for the WWW with its unique constraints and capabilities. These strategies seem analogous to a similar phenomenon that the computer industry experienced in designing online help systems, in which paper manuals were initially put online, and later approaches then evolved to design help materials specifically for the on-line medium. The authors pointed to critical differences between designing for paper and for the WWW which Web designers should keep in mind. Issues also raised in other chapters in this book were approached from this perspective, including navigation, place awareness, and the various impacts of screen resolution.

In chapter 16, "The 10 Golden Rules for Providing Video over the Web," Johnson has offered guidelines for the effective use of video on Web sites. Like Wiebe and Howe in Chapter 17, who have discussed using graphics on the

WWW, the authors of this chapter point to the potential power and the common pitfalls of incorporating video on Web sites and offers guideline for effective use. In chapter 7, Laux wrote of designing for handicapped users, and in this chapter, Johnson has raised issues of what might be called designing for handicapped hardware. Many users have speed-challenged or multimedia-challenged workstations, and Johnson has noted that designers must consider the capabilities of the average user's platform before including video. Response time on the WWW can be reminiscent of response time on large, time-shared mainframe systems of 20 years ago. As elsewhere in this book, Johnson has applied well-established general UI design guidelines to unique aspects of Web design and development. For example, Johnson pointed out that multimedia resources, like any new technology such as color, windows, or mice, should be employed only if and when their "value added" over other technologies outweighs their relative costs (in terms of download time). Johnson also employed traditional usability engineering techniques such as task analysis and usability testing to support his guidelines.

Finally, Part V explores the issues of collaboration and visualization on the WWW. In chapter 17, "Graphics Design on the Web," Wiebe and Howe, like Grose et al. in chapter 10, looked at available design guidance for Web developers, but they focused on the use of graphics in Web design, rather than on general usability issues. They noted, as did Grose et al., that most available guidance on the use of graphics on the WWW does not focus directly on usability issues, but on aesthetics and computational efficiency. They discussed usability issues relating to the use of graphic elements such as color, texture, proximity, three dimensionality, spatial layout and grouping, animation, and blinking. Guidance on how to effectively use these graphic elements appears in almost any traditional usability text, and Wiebe and Howe have applied this guidance specifically to Web design: graphics, if used effectively from a usability view point, can be a powerful way to communicate information structure and to help users build an accurate and useful user model of an information space. Although supporting efficient and effective navigation is important in any software application, it is particularly important in Web site design; the concept of a good *conceptual model* to support navigation is as crucial in Web site design as in the design of a UI for a traditional software application. Again, it is obvious that usability issues are common to both traditional software and Web design, and traditional usability principles and guidelines can be directly applied to Web design.

Chapter 18, "Collaborative Interfaces for the Web," by Greenberg, brings the field known as computer supported cooperative work (CSCW) to bear on the design of emerging software tools intended to promote and support collaborative work over the Internet. Greenberg has cited literature reporting studies of traditional collaborative work in shared work spaces (i.e., meetings) and has applied the findings to the design of a collaborative Web browser that allows many people in different locations to view, navigate, and discuss the same Web site. The literature cited in this chapter provides insights into how people in a

shared work space collaborate, how they mediate the collaboration, how they communicate (including non verbal communications such as gesturing and directing attention), how they draw attention, how they manage turn taking, how they annotate shared work artifacts such as sketches and documents, and how they monitor the group activity so as to most effectively participate in it. Greenberg and his colleagues extracted general requirements from this literature and applied them to the design of a collaborative Web browser. The initial design was usability tested and then redesigned in response to problems identified in the original design.

To me, the most interesting problem identified in the original design had to do with managing the shared view of a Web page. Maximum flexibility was built in so that all users could resize and scroll their views of the page independently. This seemed like a good idea at first, but in usability testing, it was revealed that this flexibility actually got in the way of the collaborative effort. Users could not easily tell at any moment what other users were looking at, and communicating about the common content was difficult. In the redesign, users were constrained so that they were always looking at the same view of the site and could not scroll independently of each another. This redesign eliminated complexity in the interface and facilitated communication about the content. This tradeoff between flexibility/power and simplicity in the UI occurs often in traditional UI design; most designers have a tendency to err on the side of power and flexibility and do not recognize the utility of constraints and simplicity. This example, in the context of Web design, shows the importance of balancing simplicity and ease of learning and use with functional power and flexibility.

Chapter 19, "A Zooming Web Browser," by Bederson et al., addresses the issue of designing for effective navigation and of helping users maintain a sense of place, part of a good conceptual model. The chapter describes the development of a Web browser that allows users to see the context of their navigational history and to zoom in and out on the scope of this context. This aim is accomplished by having the overall browser window present not only the current Web page, but also representations of pages from which a user navigated to the current page. Users can zoom in and out to show more or less of the pathway they have traveled to the current page (the maximum being all pages traveled in the current session and the minimum being the current page, as in traditional browsers); they can zoom in and out to show pages in the pathway in different scales (from full size to minimal detail). When users click on a link and navigate to a new page, the page is added to the growing structure representing navigational history in an animated way that helps to communicate to users how the information space they are traveling is represented by the browser. Page size in the navigational structure dynamically changes to communicate a particular page's relative distance from the page currently in focus: Pages from which users navigated to the current page become increasingly smaller as users navigate forward through successive pages, and animation helps communicate to users how the overall browser display represents their navigational actions. Bederson et al. described an iterative process of testing and redesign to refine their zooming browser: they reported in their final

test that after user training, performance on their zooming browser was superior to performance on the Netscape browser, which does not present a visual navigational history. Again, this chapter shows traditional usability engineering methods and design principles successfully applied to improve the usability of Web interfaces.

SUMMARY

The well-established field of software usability engineering, or software human factors, *could* be directly applied to the design of Web sites, but unfortunately, as in the early days of traditional software development, it is not yet being routinely applied during Web site development. Emerging lingo among Web users indicates the general lack of usability of current Web sites. The following examples were published in *Wired* and have been anonymously passed around via E-mail:

Dorito syndrome: Feelings of emptiness and dissatisfaction triggered by addictive substances that lack nutritional content. "I just spent 6 hours surfing the WWW and now I've got a bad case of Dorito syndrome."

Crapplet: A badly written or profoundly useless Java applet. "I just wasted 30 minutes downloading this stinkin' crapplet!"

World wide wait: The real meaning of WWW.

Link rot: The process by which links on a Web page become obsolete as the sites they are connected to change location or die.

Cobweb site: A Web site that has not been updated for a long time; a dead Web page.

In a recent article about the Internet, Kraut (1996) wrote that usability once on the WWW is one issue, but another issue is the usability of the process of getting connected to the Internet in the first place: "Connecting to the Internet is still an overwhelming hurdle for many families. As many online service providers are discovering, walking naive users through the vagaries of modems and networking software is expensive, but necessary" (p. 35).

I remember only too well my own experience. I resisted getting on the Internet for years, because, in spite of cheery reassurances from technically inclined clients, colleagues, and friends, I knew that it was not going to be easy, and I am not one to put up with unusable technology. Of course I was right. In spite of being computer literate, I needed 1 month to get connected. First, I had to purchase various pieces of critical software. Then, when more software was required and supplied by my access provider, I could not successfully download it, and my provider could not explain why. When I finally got the software installed, the provider sent me a long manual specifying a set of meaningless numbers that had to be entered correctly in 20 or 30 dialog box fields with meaningless labels. I followed the instructions scrupulously, and things still did not work. Again my provider seemed just as mystified as I was. Finally, after many phone calls, wasted hours, and frustration, not to mention outrage at the lack of usability of the connection process, I finally got on. Now, of course, I

cannot live or do business without the Internet. But I felt a disturbing sense of deja vu throughout the whole process of getting connected, as if I had been transported back about 20 years ago in terms of the usability of the technology. Clearly the potential of the Internet and the WWW will never be realized until the well-established principles and methods of software usability engineering are systematically applied to this new technology.

We hope that this book finds its way into the hands of the right people. We hope that usability engineers find this book helpful to understand how their expertise can be applied to Web design and that they are inspired to get involved in any Web development in their client or employment organizations. We hope that Web designers of all kinds find the book a good introduction to the usability issues involved in Web design and an inspiration to pursue other usability engineering resources to support their efforts. We hope that it does not take another 20 years for the Web development community to catch up to the traditional development community in terms of the application of usability engineering principles and methods to product development, and we hope that his book contributes to the process.

REFERENCES

Foster, R. W. (1996, October). The internet: A perspective for internetworking. *Inside DPMA*, *1*, 3.

Horton, W., Taylor, L., Ignacio, A., & Hoft, N. L. (1996). *Web page design cookbook* New York: Wiley.

Kraut, R. (1996). The Internet@Home. *ACM Communications*, *39*(12), 32–35.

Nielsen, J. (1996). Jakob Nielsen on web site usability. *Eye for Design*, User Interface Engineering), *3*(4), 1, 6–7.

Part I: PERSPECTIVES FROM PSYCHOLOGY

2 The Use of Investigatory Responses as a Measure of Learning and Memory

Charles P. Seltzer
Essex Community College

Historically, psychological research has emphasized understanding how organisms learn; during the 20th century prominent psychologists dedicated their research careers to this question. Although many noteworthy findings have been reported, a comprehensive understanding is clearly lacking (see reviews by Domjan, 1987; Spear, Miller, & Jagielo, 1990); partly because learning is often assessed and studied through behavior paradigms that use task-driven learning models. In these studies researchers have usually given an animal (often a rat, pigeon, or chick) a task to learn; the learned response is a discrete and specific behavior, such as bar pressing or maze learning. Researchers have then assessed outcome through objective measures such as maze-running speed or numbers of errors or trials to criterion. In most cases, workers directed little or no attention toward spontaneous aspects of an organism's behavior. Spontaneous exploration is a defining feature of human behavior; it can be expressed in many ways, including numerous possible interactions created by today's easy access to the Internet. As with other organisms, much human knowledge is generated through these spontaneous investigations of our world (Renner & Seltzer, 1995).

Workers have attempted to generate definitions of investigatory (exploratory) behavior, and although many different definitions have been offered, a few common themes have emerged. These common themes include the presence of spontaneous, apparently uncoerced behaviors (Barnett, 1967; Birke & Archer, 1983; Renner, 1990); independence from traditional physiological needs such as thirst and hunger (Barnett, 1975; Berlyne, 1963; Russel, 1983; Voss & Keller, 1983); behaviors whose primary purpose is to vary stimuli currently impinging on an organism (Barnett, 1967; Halliday, 1968; Russel, 1983); and behaviors that primarily allow an organism to gain information about its world (Archer & Birke, 1983; Inglis, 1983; O'Keefe & Nadel, 1979; Renner, 1988; Toates, 1983).

A recent seminal development has been the increased accessibility of computer technology. Through the Internet, people now have access to an almost limitless source of novelty and stimulation. Exploration of the Internet can meet all the needs that environmental investigation serves in other organisms: It is independent from physical needs such as hunger and thirst, it serves the primary purpose of varying current stimuli, and it allows people to gain information about the world. In this chapter I discuss the benefits associated with investigatory behaviors. Consistent with most research in this area, I emphasize animal

exploration; and the same principles apply to human exploration as well. When people investigate the resources available through the Internet, they engage in many of the same processes that animals engage in when they explore their environment. I argue that exploring the Internet is a learning experience that helps satisfy the desire for exploration. Clearly the Internet is not the only way to meet these needs: Humans have been functioning very well without the Internet for millennia. The Internet does, however, provide a new, dynamic way to express natural needs and tendencies for exposure to novelty and information.

THE RELATIONSHIP BETWEEN INVESTIGATION AND LEARNING

In the literature on investigation and exploration, a frequently mentioned function of these behaviors is to provide organisms with information about their world (Birke & Archer, 1983; Inglis, 1983; O'Keefe & Nadel, 1979; Renner, 1990; Russel, 1983). Investigatory behaviors have been thought to allow organisms to learn the locations of different resource sites (Krebes, 1978; Russel, 1983), to enhance the ability to avoid predation (Barnett, 1975; Metzgar, 1967; Renner, 1988), to detection change in familiar environments (Cowan, 1983; Russel, 1983), to allow for a detailed knowledge of home range (Shillito, 1963), and to help form cognitive maps or internal representations of an environment (Inglis, 1983; O'Keefe & Nadel, 1978; Toates, 1983). Although many of these hypotheses make sense, they have yet to be substantiated. The limited empirical support sheds little light on exactly what is learned during investigation, but indicates that the characteristics of investigatory behaviors are consistent with accepted definitions of learning.

Empirical Support

Learning has traditionally been defined as a relatively permanent or enduring change in behavior caused by environmental interaction (Domjan & Burkhard, 1982; Smith, Sarason, & Sarason, 1982). The term relatively permanent, means that the change persists after the end of the training procedure or immediate circumstances, but this change need not last forever (Domjan & Burkhard, 1982). Following periods of nonexposure, forgetting, a reduction or reversal of the behavioral change, occurs On the basis of this definition of learning, known features of investigatory behavior are consistent with other forms of learning.

Environmental investigation, involves interactions with the environment, and investigatory responses have been found to alter an organism's behavior in predictable and nonrandom manners (Inglis, 1983; Renner & Seltzer, 1995). As previously noted, when an organism is presented with unfamiliar information or resources for the first time, the usual response is one of cautious investigation. On repeated presentations, however, habituation takes place. A reasonable, learning-based explanation for this reduction in investigation is that after a period of contact with an unchanging environment, an organism acquires enough information to permit future interactions to result in diminished returns. Thus, the

amount of newly acquired information (or the amount of sensory stimulation resulting from newly acquired information) is no longer sufficient for an organism to continue to engage in investigation. In addition, some limited data indicate that if an organism is prevented from interacting with a preferred source, the intensity of the investigatory behaviors resembles that present on early contact days (Schafer, 1975).

Evidence for the Acquisition of Functional Information

Among animals, much evidence supports the relation between investigation and acquisition of functional information. For example, research has shown that the opportunity to investigate an area increases an animal's ability to escape predation (Ambrose, 1972; Metzgar, 1967; Renner, 1988). Ambrose (1972) found that allowing voles (Microtus pennsylvanicus) to investigate a barn before exposure to a barn owl reduced the threat of predation. Similarly, Metzgar (1967) discovered that pre-exposing one member of a pair of white-footed mice (*Peromyscus leucopus*) to a room reduced the pre-exposed subject's rate of predation by a screech owl, as compared to the control mouse. A more direct test of the assumption that investigation allows the acquisition of functional information, which helps an animal to escape predation, was carried out by Renner (1988). Rats that were given a chance to investigate an arena were found to escape from a mechanical simulated predator through a hole in the floor of the arena quicker than did naive subjects. By looking directly at escape behavior, this study provided evidence that investigation can increase the proficiency of adaptive behaviors.

During the investigation of an environment, memory processes aid in the forming memories of the spatial layout of the environment, including the location and attributes of stimulus objects (O'Keefe & Nadel, 1979). Several lines of evidence have been cited as support for this proposition. If a familiar object (one that has been extensively investigated) is removed from a familiar locale and placed into a novel locale, in the same arena, or in a different arena with a similar configuration, the animal behaves as if the object is no longer familiar (Poucet, Durup, & Thinus-Blanc, 1988). If the familiar object is removed and not replaced, however, an animal behaves in ways that are best interpreted as search; it spends more time in the area where the object was previously located than would be predicted by random movement (Renner & Seltzer, 1991). Studies have also shown that some organisms can use spatial information encountered during investigation to navigate around a barrier blocking a familiar part of a maze and that when organisms are introduced to a large novel environment, they establish a base or home area from which further spatial investigation originates (Menzel, 1978).

TOWARD A THEORY OF EXPLORATION AND THE INTERNET

The theory of investigation proposed here is simple; the primary supposition is that organisms have an inherent need to process sensory information. When the appropriate amount of information is not being processed, a state of disequilibrium results. Organisms regulate the amount of information being processed through a form of behavioral homeostasis, which returns the organism to a state of equilibrium. When the amount of information being processed is too low, the resulting attempt to maintain equilibrium is referred to as investigatory behavior. When the amount of information being processed is too high, organisms reduce the level of input; this behavior is referred to as withdrawal or avoidance. An additional assumption of this theory is that learning is an inherent by-product of sensory processing and therefore, an inevitable result of the behavioral homeostatic process.

A main assumption of the theory is that organisms have an inherent need, to process and assimilate information. Several lines of research support this contention. When organisms are denied sensory stimulation during very young ages, irreversible neurological and behavioral effects have often been noted. These findings appear to generalize across modalities, with long-standing neurological alterations being noted after visual deprivation (Lawicka, 1989; Lepoivre, Przybyslawski, & Bobo, 1988; Spear, McCall, & Tumosa, 1989), olfactory deprivation (Gomez-Pinilla, Guthrie, Leon, & Nieto, 1989), tactile deprivation (Vince, Lynch, Mottershead, & Green, 1987), as well as auditory deprivation (Thielcke & Zimmer, 1986). Although it is beyond the scope of this chapter to detail all studies of early sensory deprivation, taken as a whole, the studies indicate that the opportunity to process sensory information is crucial to normal neurological and behavioral development.

A vast literature maintains that if organisms are raised in captive environments (environments without sources of stimulation found in naturalistic settings), they often behave in ways best defined as maladaptive (see reviews by Erwin & Deni, 1979; Morris, 1964). Examples of maladaptive behaviors include odd limb movements, self-directed and other-directed aggressive behavior, self-directed clasping, stereotyped motor acts such as pacing and rocking, appetitive disorders, hypersexuality, as well as maternal and paternal deficits. These deficits are believed to occur because an animal is isolated from normal sources of stimulation that would be present in noncaptive situations (Boice, 1981). According to the present theory, these behaviors resulted because the behavioral homeostatic mechanisms that organisms rely on to regulate stimulus input have been circumvented by restricting a subject's behavioral options. Thus, stimulus input must be regulated by the use of additional sets of behaviors other than those observed in natural situations.

On the basis of this hypothesis, it can be predicted that abnormal behaviors resulting from captivity, termed zoo stupor by Goodwin (1975), would be attenuated by the inclusion of additional sources of stimulation. Several recent

studies, using several different mammalian species, have provided data supporting this contention (Markowitz & Spinelli, 1986; Paquette & Prescott, 1988). These studies have indicated that behaviors viewed as abnormal can be reduced by procedures as simple as adding stimulus objects (Champoux, Hempel, & Reinhardt, 1987; Paquette & Prescott, 1988), or by allowing the organisms opportunities to work actively for food rather than to receiving free food.

Animal enrichment studies have demonstrated that in the laboratory, direct comparisons can be made between organisms raised in enriched and impoverished captive conditions (see reviews by Greenough, 1984; Lynch & Wells, 1978; Renner & Rosenzweig, 1987). Research has shown that giving laboratory-raised rats early experiences with novel objects in enriched conditions results in numerous facilitative effects on future learning as well as nervous system development. Nervous system modification can occur after very short exposure periods to enriched conditions, with measurable neurological changes noted after only 40 minutes of enriched experience (Ferchmin & Eterovic, 1986). Enrichment effects have been noted in adult and preadult organisms, although the magnitude of the enrichment effect is usually found to be greater preadults (Bennett, Rosenzweig, Morimoto, & Hebert, 1979). Finally, it has been shown that in most measures, direct contact with the objects is required for these neural changes to occur (Ferchmin, Bennett, & Rosenzweig, 1975).

THE INTERNET AND EXPLORATION

The internet offers a plethora of opportunities to process unique and novel forms of information. Everyone has an inherent need for information and stimulation. Individuals can achieve their desired levels of stimulation in many ways and the Internet is but one avenue. What makes the Internet unique is its ability to share information and stimulation with a seemingly infinite number of individuals and sources. Although many argued for and against the Internet, learning appears to be an by-product of Internet exploration. People may also argue, however, that if Internet use does provide enough stimulation (as well as information processing) to prevent disequilibrium, it may make other avenues of investigation (e.g., reading, family interaction) less likely. People who spend much time on the internet apparently do ignore other activities. The question, however, is; do the gains associated with Internet use justify reallocating time away from traditional investigatory situations?

One way to investigate the Internet is to regard it as environment enrichment. The Internet greatly facilitates exposure to information in much the same way that adding objects to an organism's environment increases its level of information and stimulation. As with any stimulation, these exposures could be postulated to alter adaptiveness, or if nothing else, to relieve boredom. Researcher may never learn whether or not the Internet has the same neurological effects on people as an enriched environment has on rats, but the Internet can fundamentally change the way people think about the world, and the way they process information. Having information at their fingertips frees people from the

doldrums of data collection so that they can spend time using and interpreting this information.

REFERENCES

Ambrose, H. W. (1972). Effect of habitat familiarity and toe clipping on rate of owl predation on *Microtus pennsylvanicus*. *Journal of Mammology*, *53*, 909–912.

Archer, J. A., & Birke, L. I. A. (1983). *Exploration in animals and humans*. Cambridge, England: Cambridge University Press.

Barnett, S. A. (1967). Instinct and intelligence: Behavior of animals and men. Englewood Cliffs, NJ: Prentice Hall.

Barnett, S. A. (1975). *The rat: A study in behavior*. Chicago: University of Chicago Press.

Bennett, E. L., Rosenzweig, M. R., Morimoto, H., & Hebert, M. (1979). Maze training alters brain weights and cortical RNA/DNA ratios. *Behavioral and Neural Biology*, *26*, 1–22.

Berlyne, D. E. (1963). Motivational problems raised by exploratory and epistemic behavior. In S. Koch (Ed.), *Psychology: A study of a science* (pp. 288–299). New York: McGraw-Hill.

Birke, L. I., & Archer, J. (1983). Some issues and problems in the study of animal exploration. In J. A. Archer & L. I. Birke (Eds.), *Exploration in animals and humans* (pp. 1–16). New York: Von Nostrand Reinhold.

Boice, R. (1981). Captivity and feralization. *Psychological Bulletin*, *89*, 407–421.

Champoux, M., Hempel, M., & Reinhardt, V. (1987). Environmental enrichment with sticks for singly-caged adult Rhesus monkeys. *Laboratory Primate Newsletter*, *26*, 5–7.

Cowan, P. E. (1983). Exploration in small animals: Ethology and ecology. In J. A. Archer & L. I. Birke (Eds.), *Exploration in animals and humans* (pp. 147–170). New York: Von Nostrand Reinhold.

Domjan, M., & Burkhard, B. (1982). *The principles of learning and behavior*. Monterey, CA: Brooks/Cole.

Erwin, H., & Deni, R. (1979). Strangers in a strange land: Abnormal behaviors or abnormal environments. In J. Erwin, T. L. Maple, & G. Mitchell (Eds.), *Captivity and behavior* (pp. 12–58). New York: Von Nostrand Reinholt.

Ferchmin, P. A, Bennett, E. L., & Rosenzweig, M. R. (1975). Direct contact with enriched environments is required to alter cerebral weights in rats. *Journal of Comparative and Physiological Psychology*, *88*, 360–367.

Ferchmin, P. A., & Eterovic, V. A. (1986). Forty minutes of experience increases the weight and RNA content of cerebral cortex in periadolescent rats. *Developmental Psychobiology*, *19*, 511–519.

Gomez-Pinnilla, F., Guthrie, K. M., Leon., M., & Nieto, S. M. (1989). Receptor increase in the olfactory bulb of the rat after early odor deprivation. *Developmental Brain Research*, *48*, 161–165.

Goodwin, L. G. (1975). Scientific work of the zoological society of London. In *Research in zoos and aquariums* (pp. 35–50). Washington, DC: National Academy of Sciences.

Greenough, W. T. (1984). Possible substrates of plastic neural phenomenon. In G. Lynch, J. L. McGaugh, & N. M. Weinberger (Eds.), *Neurobiology of learning and memory* (pp. 470–478). New York: Guilford.

Halliday, M. S. (1968). Exploratory behavior. In L. Weiskrantz (Ed.), *Analysis of behavioral change*. New York: Harper & Row.

Inglis, I. R. (1983). Towards a cognitive theory of exploratory behaviour. In J. A. Archer & L. I. Birke (Eds.), *Exploration in animals and humans* (pp. 73–103). New York: Von Nostrand Reinhold.

Krebes, J. R. (1978). Optimal foraging: Decision rules for predators. In J. R. Krebs & N. B. Davies (Eds.), *Behavioral ecology: An evolutionary approach* (pp. 23–63). Oxford, England: Blackwell.

Lawicka, W. (1989). Delayed response to light stimuli in binocularly deprived cats. *Acta Neurobiologia Experimentalis*, *49*, 73–92.

Lepoivre, M., Przybyslawske, J., & Bobo, G. (1988). Visual field in dark reared cats after an extensive recovery period. *Behavioral Brain Research*, *38*, 245–251.

Lynch, G., & Wells, J. (1978). Neuroanatomical plasticity and behavioral adaptability. In T. Teyler (Ed.), *Brain and learning* (pp. 105–124). Stamford, CA: Greylock.

Markowitz, H. (1982). *Behavioral enrichment in the zoo*. New York: Von Nostrand Reinhold.

Markowitz, H., & Spinelli, J. (1986). Environmental engineering for primates. In K. Benirschke (Ed.), *Primates: The road to self-sustaining populations* (pp. 489–498). New York: Springer-Verlag.

Menzel, E. W. (1978). Cognitive mapping in chimpanzees. In S. H. Hulse, H. Fowler, & W. K. Honig (Eds.), *Cognitive processes in animal behavior*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Metzgar, L. H. (1967). An experimental comparison of screech owl predation on resident and white footed mice (*Peromyscus leucopus*). *Journal of Mammology*, *48*, 387–391.

Morris, D. (1964). The response of animals to a restricted environment. Symposium of *Zoological Society of London*, *13*, 99–108.

O'Keefe, J., & Nadel, L. (1979). Precis of O' Keefe and Nadel's The hippocampus as a cognitive map. *Behavioral and Brain Sciences*, *2*, 487–533.

Paquette, D., & Prescott, J. (1988). Use of novel objects to enhance environments of captive chimpanzees. *Zoo Biology*, *7*, 15–23.

Poucet, B., Durup, M., & Thinus-Blanc, C. (1988). Short-term and long-term habituation of exploration in rats, hamsters and gerbils. *Behavioral Processes*, *16*, 203–211.

Renner, M. J., (1988). Learning during exploration: Evidence for acquisition of functionally significant knowledge during spontaneous exploratory activity. *International Journal of Comparative Psychology*, *2*, 43–56.

Renner, M. J. (1990). Neglected aspects of exploratory and investigatory behavior. *Psychobiology*, *18*, 16–22.

Renner, M. J., & Rosenzweig, M. R. (1987). *Enriched and impoverished environments: Effects on brain and behavior*. New York: Springer-Verlag.

Renner, M. J., & Seltzer, C. P. (1991). *Behavioral grammars used in object investigation by Long-Evans rats*. Poster presented at the annual Meeting of the American Psychological Society. Washington, DC.

Renner, M. J., & Seltzer, C. P. (1995). Sequential structure in behavioral components of object investigations by Long-Evans rats. *Journal of Comparative Psychology*, *108*, 71–81.

Russel, P. A. (1983). Psychological studies of exploration in animals: A reappraisal. In J. A. Archer & L. I. Birke (Eds.), *Exploration in animals and humans* (pp. 1–16). New York: Von Nostrand Reinhold.

Schafer, A. B. (1975). Newborn responses to nonsignal auditory stimuli: Electro-encephalographic desynchronization. *Biophysiology, 12*, 359–366.

Shillito, E. E. (1963). Exploration in the short-tailed mole (*Microtus agrestis*). *Behaviour, 21*, 145–154.

Smith, R. E., Sarason, G., & Sarason, B. R. (1982). *Psychology: The frontiers of behavior*. Cambridge, MA: Harper & Row.

Spear, N. E., Miller, J. S., & Jagielo, J. A. (1990). Animal memory and learning. *Annual Review of Psychology, 41*, 169–211.

Spear, P. D., McCall, M. A., & Tomaso, N. (1989). W and Y cells in the C layers of the cat's lateral geniculate nucleus: Normal properties and effects of monocular deprivation. *Journal of Neurophysiology, 61*, 58–73.

Thielcke, G., & Zimmer, U. (1986). Early experience determines the song of the chiffchaff. *Ethology, 73*, 191–196.

Toates, F. M (1983). Exploration, a motivational and learning theory: A cognitive incentive view. In J. A. Archer & L. I. Birke (Eds.), *Exploration in animals and humans* (pp. 1–16). New York: Von Nostrand Reinhold.

Vince, M. A., Lynch, J. J., Mottershead, B. E., & Green, G. C. (1987). Interactions between normal ewes and newly born lambs deprived of visual, olfactory and tactile sensory information. *Applied Animal Behaviour Science, 19*, 119–136.

Voss, H. G. & Keller, H. (1983). *Curiosity and exploration: Theories and results*. San Diego, CA: Academic Press.

3 Visual Information Processing on the World Wide Web

William Marks
The University of Memphis

Cynthia L. Dulaney
Xavier University

Information and knowledge about the physical environment are acquired mainly through visual perception. The physical environment contains surfaces, objects, and events that perceivers use to adapt to and make sense of their environment. The World Wide Web (WWW) is an electronic, hypermedia environment that contains information which is present as text, images, sounds, and apparent motion displays signaling object and event transformation or change. Three-dimensional virtual objects and environments using virtual reality modeling language (VRML) now permit users to interactively perform actions and to examine scene content from multiple vantage points (Nadeau,1996). Because the WWW can be explored and the information provided by its environment can be sampled by users for purposes of discovery, it is appropriate to consider issues in visual information processing relevant to interactions with this environment. This chapter provides a highly selective overview of several topics about the visual processing of information that are pertinent to an appreciation of perceivers in the development of Web environments. We review early, low-level, visual processes that provide information to the perceptual system, as well as higher level processes that result in conscious awareness of the external world and that permit actions to be taken on the external world. Our goal is to present principles of visual information processing relevant to the design of Web pages ; such priniciples may aid designers to develop Web pages that are easy to search and browse and that promote interaction.

DISTAL AND PROXIMAL STIMULI

Perception typically provides veridical information about the physical environment because human sensory-perceptual systems have evolved to detect features that specify how the physical world is structured from a particular vantage point and at a particular time. Over time, perceptual systems detect the

environment's invariant properties that uniquely specify objects and events independent of any spatial-temporal constraints (Gibson, 1979). Nevertheless, information derived from sensory systems can sometimes result in perceptual errors as when people perceive objects and events that are not based on the true state of affairs of the physical world. In these cases, perception and reality do not correspond, and people perceive illusions. These perceptual illusions rarely occur, but they are remarkably revealing about how perceptual systems make assumptions about the external environment on the basis of the evidence provided by sensory systems and perceivers' previous experiences (Gillam, 1980).

The fact of perceptual illusions highlights a distinction between distal and proximal stimuli. Distal stimuli refer to objects and events in the physical world, whereas proximal stimuli refer to representations of distal stimuli projected on the retina, which contains light-sensitive receptors necessary for perceiving. An external-world distal stimulus is accessible to perceivers only by virtue of a proximal stimulus. The proximal stimulus represents what in the world the eventual percept is about. Proximal stimuli, however, are only two-dimensional projections of distal stimuli and proximal stimuli constantly change in response to, for example, distance from a distal stimulus (i.e., a proximal stimulus becomes larger or smaller). A proximal stimulus also moves across the retina as the head or eyes move. With such variability in proximal stimuli, it is surprising that perceivers can detect the constant properties of distal stimuli, and that perceptual illusions are infrequent.

DATA-DRIVEN PROCESSING AND CONCEPTUALLY DRIVEN PROCESSING IN PERCEPTION

Bottom-up, data-driven processing must be distinguished from top-down, conceptually driven processing in perception. Data-driven processing refers to the process of extracting information about the characteristics and properties of a proximal stimulus, such as color, shape, orientation, and size, to be sampled by the perceptual system. These data trigger responses in separate visual processing streams that are sensitive to form, color, depth, and motion (see, e.g., Maunsell & Newsome, 1987; Van Essen, 1985). The modularity of perceptual processing, inwhich many partially autonomous processors rapidly perform limited and independent functions, is now a widely accepted view (Fodor, 1983; Kosslyn & Koenig, 1992; Marr, 1982). Bottom-up, data-driven processing occurs because these low-level processing modules are able to independently compute features of proximal stimuli to which they are sensitive. One familiar example of how perceivers may arrive at integrated percepts by computations in data-driven processing modules are random-dot stereograms developed by Bela Julesz (1971). In such displays, two separate fields of uniform, randomly generated textures of black and white dots form two disparate images because the central area in each image has been slightly displaced in opposite directions to left and right. Neither image, however, contains a recognizable form, nor are any cues in the displays

available to produce the apparent depth effect that emerges, of a square floating on a background, when the two fields are viewed in a stereoscope.

The conclusion drawn from random-dot stereograms is that low-level visual information processing can proceed in a data-driven fashion to provide reliable information for perceiving depth based on the disparity between two images. This correspondence problem in stereopsis is the ability to judge depth binocularly because of binocular disparity (the fact that each eye registers a slightly different image). It is assumed that there are processing modules sensitive to corresponding and noncorresponding points in the neural processing of a proximal stimulus. Corresponding points are fused into a single image, but noncorresponding points cannot be fused and therefore serve as information for stereopsis. Thus, data-driven processing enables the extraction of information from a proximal stimulus without the necessity of previous knowledge.

The possibility that data-driven processing results in reliable percepts also highlights the fact that distal stimuli and their proximal counterparts are typically rich sources of information about the layout of the environment and the objects and events in the environment (Gibson, 1979; Marr, 1982). Nevertheless, the data provided by sensory systems may sometimes result in incorrect percepts about the physical world because the data may be limited with respect to object and event specification. Camouflage is a good example of how data-driven processing can result in percepts that do not provide accurate information about the external world (see Hailman, 1977). Other examples of how data-driven processing can result in incorrect percepts are discussed in several of the following sections.

Bottom-up, data-driven processing can be contrasted with top-down, conceptually driven processing, which emphasizes the importance of perceivers' concepts in determining perception. In general, perception refers to the interpretation of sensory events so that the data from the events are given meaning and organization. Perceptions are conscious experiences of objects and events in the world. Perceptions about apples are not limited to their sensory properties such as bright, red, and round (although these may be perceptual descriptions of form; see Rock, 1983). Therefore, a unifying principle in research on visual perception has been that perception also requires that conceptually-driven processes be brought to bear on sensory data. During perceptual processing, perceivers make use of previous knowledge to interpret incoming sensory data. They make inferences and hypotheses and also generate expectations about the physical world that give rise to a particular proximal stimulus. This process of perceptual inferencing further constrains and guides how incoming sensory data are interpreted. In most instances, the conceptual framework that perceivers make use of is highly situational, or context dependent. Perceivers expect that the tall rectangular shape in a kitchen is a refrigerator rather than a shower stall. The percept is that of a refrigerator unless sensory data provide unambiguous information about the properties that would specify a shower stall.

Thus, perceptual errors or illusions occur when the perceptual system misinterprets incoming sensory data because the context in which the data are presented produces an expectation that the data are of a particular type. Readers of

text often have the experience of anticipating that a particular word will occur in a particular context. They find themselves subvocally speaking the word although a different word is being processed by the sensory system at the time the anticipated word is spoken. We elaborate on the importance of context effects in perceptual processing in a later section, but we next discuss some of limits on sensory information processing that constrain the kinds of information available to perception.

SENSORY CONSTRAINTS ON PERCEPTION

Human visual systems are sensitive to the portion of electromagnetic radiation referred to as light. Information about an object's color, size, and shape is conveyed by light. Reflected light from objects in the world bridges the distance between perceivers and objects. In the eyes, reflected light is focused on the retina, which, as previously mentioned, contains photoreceptors that convert light energy into neural signals. The neural signals convey all the information about a distal stimulus, such as color, brightness, spatial location, motion, and depth, and therefore provide the foundation for visual perceptions. Human retinas contain two types of photoreceptors – cones and rods. Cones are differentially sensitive to wavelengths of light which when reflected from surfaces and objects provides the basis for perceiving color or hue. Cones are located mainly in the macula area of the human eye, which has at its center a pitted area called the fovea. Outside of the area of fovea, the number of cones decreases, and the number of rods increases. With respect to function, the photoreceptors in human eyes provide a "duplex solution" to two competing demands – providing good spatial resolution and detecting dim light. Cones provide perceivers with good resolution, whereas rods provide perceivers with high sensitivity. Resolution refers to the ability to distinguish differences in the spatial details of light in a proximal stimulus when light is in abundance. Sensitivity refers to the ability to detect small amounts of light when it is sparse. A rod photoreceptor is considered to be at the theoretical limit of sensitivity because it can respond to a single photon (a quantum of electromagnetic radiation).

This duplex solution constrains the aspects of proximal stimuli that observers can discriminate at any one time. A key constraint concerns visual acuity as a function of retinal location. Visual acuity is a measure of resolution and, as would be expected, it is greatest in the area of the retina that contains cones. Visual acuity decreases drastically with increasing distance from the fovea and toward the retinal periphery that contains mainly rods. This phenomena can be seen by looking at a letter on a computer screen: Without moving the eyes, only a few letters to either side can be read. Anstis (1974) provided an interesting demonstration of the difference in visual acuity between the foveal area and the retinal periphery. He computed the sizes that letter characters at various distances from the fovea must be to be detected with acuity equal to a letter projected onto the fovea. He found that smaller letters had to be brought much nearer to the

foveal area than did the larger letters. In other words, the size of the letters increased at the same rate that visual acuity decreased. Thus, the relation between visual acuity and cone density is an important constraint in the design of displays. Information that is too far from the center of focus may not produce the intended effects of having viewers notice it, or acquire much information from the displaced item unless the center of focus is shifted or there is an abrupt visual onset of information (Yantis & Jonides, 1990).

Factors other than cone density should also be considered. Visual acuity decreases appreciably as the level of illumination drops because cones are not sensitive in dim light and rods do not provide the resolution necessary for good acuity. Acuity also decreases with age. When black letters in eye charts are changed to light gray letters, which reduce contrast (by reducing the difference in light intensity between letters and background), older adults have greater difficulty reading the eye chart (Pelli, Robson, & Wilkins, 1988). Even when younger and older adults are matched for acuity on a high-contrast eye chart, older adults show a much greater decline in acuity than do younger adults on a low-contrast eye chart (Adams, Wang, Wong, & Gould, 1988). Another factor that affects acuity concerns the "crowding effect" (Anstis, 1974). People have difficulty reading displays when the letters are too close together; smaller letters on a less crowded line are easier to read than are larger letters on a crowded line. The crowding effect may partly explain why people do not like to read text that only consists of capital letters. Sanders and McCormick (1993) have provided recommendations about the size and density of alphanumeric characters on computer displays, but as the amount of potential information in a display increases, there is likely to be a reduction in a perceiver's ability to detect all the information that a designer intended to be noticed (i.e., crowding effects may occur).

In the next section, we briefly discuss the spatial frequency analysis approach. This approach has been useful in issues of visual sensitivity, and it has been applied to several levels of perceptual processing, including pattern recognition. Although it is well beyond the scope of this chapter to attempt a detailed account of the insights from the spatial frequency analysis approach, we provide a brief, superficial description of this topic because of its influence in research and theory in visual perception (see Marr, 1982; also see Sekuler & Blake, 1994, for a introduction).

Contrast Sensitivity and Spatial Frequency

Visual acuity measures only the ability to resolve small spatial details, but other spatial properties of stimuli are important in perceptual processing. These include overall size, orientation, and contrast. The contrast sensitivity function is a measure of spatial sensitivity characterizes the ability to detect fine details as well as global aspects of the stimulus. The spatial contrast sensitivity function measures contrast sensitivity to grating patterns consisting of parallel vertical or horizontal bars with varying widths and contrasts. Four separable properties of grating patterns can be measured: spatial frequency, contrast, orientation, and

spatial phase. Spatial frequency refers to the width of the bar gratings measured by the number of light and dark bars in a given distance on the retina. Orientation refers to the alignment of the bars in the grating. Contrast is related to the differences in intensity between the bars in the grating. Spatial phase refers to the position of the grating relative to some landmark (typically, beginning with the grating's left edge). Each of these components can be varied without affecting the others. For present purposes, spatial frequency is the most important of these components because it has been so strongly implicated in the perception of form and pattern.

Visual stimuli can be analyzed in terms of the spatial frequencies in the stimuli. High spatial frequencies are sharp edges and fine, spatial details, whereas low spatial frequencies are global, coarse spatial aspects of a stimulus. The different spatial information conveyed by high and low spatial frequencies can be demonstrated by viewing either a defocused or diffused stimulus. Simply blurring an image defocuses a stimulus by reducing the contrast of the high spatial frequencies but not of the low spatial frequencies. Blurring an image, as a consequence, eliminates sharp edges and fine details. Recognizing that an image is, for example, a tree, however, is not affected by stimulus blurring because coarse global properties of the image, such as the trunk, larger limbs, and clumps of leaves are unaffected by blurring, but it would be nearly impossible to see the leaves' edges, the smaller branches, or the individual blades of grass beneath the tree. Clearly, the low spatial frequencies of a stimulus convey important information for recognizing a stimulus.

The contrast of both high and low spatial frequencies is reduced if an image is diffused. Recognition of a diffused image is extremely difficult. Diffusion is captured by looking through, for example, a fogged windshield or wax paper. In each case, both high and low spatial frequencies are filtered, and stimulus recognition is greatly impaired. The differences observed between the effects of diffusion and defocusing on recognition performance points to the importance of both high and low spatial frequency information in processing visual information.

Contrast influences perception in other ways. *Lightness contrast* occurs when an object of fixed intensity appears either lighter or darker because of the lightness of its surround. For example, a gray stimulus of moderate intensity appears dark when an adjacent stimulus is bright and light when its neighbor is dark. Similarly, *color contrast* occurs when a fixed color looks different because of the color of its surround. For example, a gray stimulus appears slightly greenish against a red background, and yellow against a blue background. A classic demonstration of color contrast appears in Land (1959) who showed the importance of contrast in the experience of objects' colors. The implications of these simultaneous contrast effects for the perception of displayed information are clear: At this relatively low level of visual processing, context effects can have a tremendous impact on perception. The perceptual system generates descriptions of sensory data based on the cues that it has available (Rock, 1983). In the next section, we discuss other cues that used to perceive depth and distance. Although

many Web pages are text based and are, therefore, best presented as two-dimensional displays, designers may wish to develop three-dimensional graphic representations to enhance the quality of their presentations. Depth perception is particularly important for understanding how perceivers are able to recover three-dimensional perspectives in two-dimensional displays.

Perceiving and Depicting Depth and Distance

Perceptions of the environment are three dimensional, but only two of these dimensions, horizontal and vertical, appear in a two-dimensional representation of the environment in a proximal stimulus. Depth is not geometrically represented on the retina. Nevertheless, a two-dimensional proximal stimulus does provide some depth cues. We have already mentioned stereopsis, provided by the binocular disparity between the two eyes, as a basis for depth perception. Other depth cues are present under monocular and static viewing conditions; they are sometimes referred to as pictorial depth cues because they can be depicted on a two-dimensional surface. These pictorial depth cues, which are extremely powerful in conveying information about distance and depth, include occlusion (interposition) where one surface partially covers another surface. An occluded surface appears farther away than an occluding surface. For example, a book held in the hands is nearer to the eyes than the surface just beyond it, including the hands, because the book partially or completely occludes these surfaces.

Linear perspective provides cues to distance when physically parallel edges seem to converge with increasing distances, as occurs when viewing a long hallway. The cue of aerial perspective indicates that surfaces at a distance have lower contrast, fewer spatial details, and lower spatial frequencies. Size is another depth cue when objects farther away are seen as smaller than nearer objects. There is also the cue of shading that produces a gradient in the amount of reflected light across a three-dimensional object. This gradient provides a depth cue because shadows are not cast by flat, two-dimensional surfaces. Ramachandran (1988) has demonstrated that a compelling effect of shading is provided when a shaded area is enclosed by an outline. An illusory contour will yield an even stronger effect of shape from shading than a real outline.

Texture gradient as a cue for distance is indicated by a greater density of texture elements at increasing distances. There are gradients of change so that, for example, blades of grass appear more densely compacted at a distance than beneath the chin of a viewer. Texture gradients are such powerful depth cues that they can capture depth in simple, schematic representations of shape, even in the absence of shading information (Todd & Reichel, 1990). Discontinuities in texture gradients provide compelling information about the spatial layout of objects, including where sharp edges, ridges, and tilts occur along the surfaces of objects (Gibson, 1979). Linear perspective and texture gradient are also important cues for size constancy. Size constancy refers to the idea that an object's size will remains relatively unchanged when the object is seen at varying distances. Perceptual constancy is also maintained for lightness, color, and shape. In each case, the perceptual system seems to take into account aspects of a proximal

stimulus to achieve a stable, invariant percept despite changes in the appearance of the environment.

Our purpose in this section has been to highlight the role of static pictorial cues to depth and distance in two-dimensional displays. Flat, two-dimensional representations are often less appealing and may be less likely to capture and maintain attention than are three-dimensional representations that use pictorial depth cues. In many instances, designers of Web pages effectively use these depth cues to enhance their pages as well as to facilitate searching and browsing. In the next section we consider how the components of a display are organized in perceptual processing to yield the emergence of meaningful patterns.

ORGANIZATIONAL PRINCIPLES IN VISUAL INFORMATION PROCESSING

In processing visual displays, some aspects are perceived as figure, or that which is attended to, and other aspects are considered background, or that which is unattended to. Several principles of organization facilitate processing the components of a visual display into meaningful figures or units (see Rock & Palmer, 1990). Components of similar identity are usually perceived as belonging together and thus forming a coherent unit. Components that are close to each other are also usually perceived as forming a unit. In this case, the unitization occurs because of proximity or of similarity in location. Visual processing tends to prefer contours that continue smoothly along their original course, without abrupt change. This principle, known as good continuation is used effectively in many so-called impossible figures in which the perceptual system is sometimes slow to detect anomalies when processing the figures (Hochberg, 1970). Similarly, things that have a common fate by moving in the same direction are perceived as forming a unit. Furthermore, a symmetrical, balanced presentation of components aids in perceiving of a group of items as belonging to a coherent unit. These principles were formulated by the Gestalt psychologists who emphasized the importance of overall structure or pattern as the major deteminant of form perception (Rock & Palmer, 1990). Taking these organizational principles into account when designing Web page allows users of the page to quickly grasp the major components or units of the page and to quickly direct attention to relevant information.

Once attention is directed to a particular unit of a visual display, perceivers usually process the global aspects of a unit before processing the features of the unit (see Kimchi 1992, for a review). This initial preferential processing of global information is referred to as the *global precedence effect* (Navon, 1977). Designers of Web pages must be aware of the organization of a page in relation to the global precedence effect. Because viewers often quickly scan a page, to the extent possible, the most relevant information should be presented at a global level because this level is first perceived by viewers (Navon, 1977). At the least, it

is helpful to provide a global representation of the page's contents because of the precedence given to global structure in perceptual processing.

The global precedence effect is contingent on several factors. As mentioned earlier, the size of the unit or stimulus and the degree to which the stimulus is presented to the foveal area are important (see the section titled Sensory Constraints on Perception). Global precedence most frequently occurs when the visual angle (the angle subtended by an object on the eye) at which the stimulus is viewed is from 4 to 7 degrees (Kinchla & Wolfe, 1979). On the basis of the average viewing distance from a computer monitor, this range encompasses much information presented on Web pages, but the global precedence effect can occur at larger angles under some circumstances (Navon & Norman, 1983). Another factor that affects the global precedence effect is referred to as the level-readiness effect (Hoffman, 1980; Ward, 1982). If information has just been processed at a global level, then a global precedence effect is more likely. But, if information has just been processed at a featural, or local, level, then subsequent information may be processed more readily at this level. Web designers should consider these aspects of global processing when linking several pages together. For example, if a user connects from one page to another, a designer should keep in mind the level at which a user is most likely to process information to facilitate subsequent processing.

Pattern Recognition

Form perception occurs when a figure's components are organized into a coherent percept, but,perceivers have had previous experiences with many of the forms that they encounter. Previous experience with a form may result in a conscious awareness of familiarity and recognition. Pattern recognition refers to the ability to classify patterns as instances of a category based on stored memory representations of perceived patterns. Pattern recognition occurs when outputs from earlier stimulus-driven processes provide a pattern of representation that can be used by conceptually driven processes to identify a pattern. As we have previously pointed out, stimulus-driven and conceptually driven processing often occur simultaneously and interactively. This interactive bottom-up and top-down processing limits the sensory inputs that are relevant to pattern recognition and promote rapid classification of a pattern. Various patterns that can be recognized by human perceivers include letters, words, faces, objects, scenes, and events. Perhaps more impressive is perceivers' ability to recognize separate instances of each of these categories (e.g., different dogs), as well as to detect novel instances in categories (new breeds of dog) and transformations of instances (the same dog standing or lying down). In this section, we briefly discuss a recent theory of object recognition, referred to as recognition by components theory (Biederman, 1987), that accounts for the success of pattern recognition in humans. We also discuss the role of conceptually driven processes as provided by context effects in object recognition. We do not discuss theories of word recognition, but instead limit our discussion of word recognition to how words are perceived and recognized in various contexts (see Carr & Pollatsek, 1985, for a review of

theories of word recognition; see Pinker, 1984, for a review of theories of object recognition).

According to Biederman's theory when a perceiver views a visible object the visual system first creates a rough classification of the object on the basis of three-dimensional shape. This level of classification is the basic-level category typically used to identify objects as, for example, a chair (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Perceivers do not typically identify objects at the superordinate level (furniture) because this level is too general; they do not identify objects at the subordinate-level category (office chair) because this level is too specific for most purposes. Biederman argued that because the initial classification occurs at the basic level, details needed for fine discriminations are not used in the initial stages of object recognition. Biederman proposed that any view of an object can be represented as the relation between a few simple three-dimensional shapes called *geons* (geometric ions). The geons are a set of approximately 24 simple volumes (e.g., cylinders, cones, blocks) that serve as the basic building blocks of all objects that can be recognized, just as 26 letters serve to encode all written English. Pattern recognition consists of specifying which geons are present and describing the relations among the geons. Biederman has suggested that only about three geons are necessary to describe most objects. Thus, one virtue of Biederman's theory is its ability to represent a very large number of possible objects by means of a few components.

Other aspects of the theory make it appealing as well. For example, the geons can be distinguished from one another from almost any viewing perspective; and they can typically be recognized even when randomly chosen parts of any geon is occluded or eliminated. Although objects are recognized faster if they are viewed from their most typical perspective (Palmer, Rosch, & Chase, 1981), recognition by components theory can explain the difficulty of recognizing an object seen in an unusual perspective. A blender is difficult to recognize if it is viewed from directly overhead because the proximal stimulus provides no hint of the geons (the handle and the base) that are critical for identifying the item as a member of the category *blender*. Three-dimensional object representations are useful for providing viewpoint-independent displays that are easily identified. Such displays may be particularly effective in VRML environments because users can recognize objects in the environment despite the changing perspectives that their movements through the environment produces, but it is important to keep in mind that some object perspectives are more difficult to recognize than are others when units used for recognition are obscured.

The Role of Scene Context in Object Recognition

In Biederman's theory, once the appropriate geons have been extracted from a proximal stimulus and the relations among the geons have been established, the object is identified from matches against stored memory representations of object models, but object processing typically does not occur in isolation. Objects are often perceived in larger scene contexts that provide world-knowledge, representational frameworks (scene schemata) for perceiving objects. Although

these top-down, conceptually driven processes can facilitate object recognition, they can also hinder it. Palmer (1975) showed that the identification of ambiguous objects is facilitated when the objects are viewed in an appropriate scene context (a loaf of bread in a kitchen), but not in an inappropriate scene context (a mailbox in a kitchen). Object identification was slower when viewed in an inappropriate scene context than when viewed in isolation.

Perhaps part of the reason for the slowed recognition of unexpected objects in scene contexts is due to conscious attention being drawn to the inappropriate item in the scene (Loftus & Mackworth, 1978). Scene recognition occurs almost automatically as long as certain constraints are not violated (Biederman, 1983). The probability of an object being in a scene is one such constraint; there is a low probability that an octopus will appear in a farm scene, and therefore the octopus is unexpected. Unexpected objects are focused on first and for a longer duration than are expected objects (Loftus & Mackworth, 1978), a process that disrupts normally fast scene recognition. When a natural scene is presented for as little as 100 milliseconds, people are able to describe the overall nature of the scene and to identify nearly half of the objects present (Biederman, Glass, & Stacy, 1973). Biederman (1990) suggested that geons and geon clusters may explain the ease with which complex scenes are recognized. Apparently, perceivers are able to use the normal relations in a scene to quickly identify the scene and its components. The normal relations provide a context that influences the perception of the scene. On the other hand, scene recognition and object identification are slowed when the normal relations in a scene are disrupted by randomly rearranging the objects (Biederman, Rabinowitz, Glass, & Stacy, 1974).

The results of these studies indicate the importance of context in object recognition. Perceivers are able to use their world-knowledge in a top-down, conceptually driven fashion to easily recognize scenes and their components. Context provides a representation for viewing a scene display, but, if the scene or object is inappropriate, then scene recognition and object processing are delayed. There are, however, some advantages to viewing either unexpected objects or scenes because these items capture attention for longer duration and are better discriminated on a later occasion than are expected items (Pezdek, Whetstone, Reynolds, Askari, & Dougherty, 1989). It is important not to confuse unexpected items (e.g., a fire hydrant in an office scene) with irrelevant items (e.g., a bulletin board). Unexpected items capture attention and are discriminated better than are expected or irrelevant items.

The Role of Context in Word Recognition

Words, like objects, seldom occur in isolation. Context effects in the perception and recognition of words have been studied in cognitive psychology primarily from the viewpoint of how conceptually driven processes facilitate encoding words and accessing their meaning. Many studies have been conducted using a lexical decision task, in which subjects must decide whether letter strings are words or nonwords (e.g., Meyer & Schvaneveldt, 1974). People can more quickly decide that an letter string (e.g., DOCTOR) is a word if it is preceded by a related

word (e.g., NURSE) rather than a unrelated word (e.g., BREAD). These results are typically described in terms of semantic priming of linked concepts by a process of spreading of activation from the initial word's representation in memory to words linked to the concept. Unrelated words are less likely to be primed by the spreading activation from the initial concept because the activation grows weaker as it spreads to more remote concepts (for a review, see Neely, 1991). Research evidence for the role of conceptually driven processing in producing these semantic priming effects has demonstrated that the magnitude of priming is affected by how the initial items are processed (e.g., Friedrich, Henik, & Tzelgov, 1991). There is virtually no priming effect if an initial word is judged for its physical attributes (e.g., uppercase) before the presentation of the related word, and there may even be inhibitory effects under these circumstances (Besner, Smith, & MacLeod, 1990). Similar priming effects occur when words are presented as completions for sentences (e.g., Tulving & Gold, 1963). Much like the scene context effects mentioned earlier, a word in context is easier to recognize than a word in isolation. The effect is also obtained with appropriate contexts, but not with ill-formed or inappropriate contexts (Potter, Moryadas, Abrams, & Noel, 1993).

The idea of linked connections between nodes in a network is an important aspect of the hypertext-based information presentation system WWW. Interacting with hypertext may require complex visual and conceptual representations of the structure of information organization (Chen & Rada, 1996). Designers of hypertext-based Web pages may find it instructive to appreciate the way in which nodes in a network can be linked to facilitate the organization of knowledge (see, Fisher, Faletti, Patterson, Thornton, Lipson, & Spring, 1990; Holley & Dansereau, 1984; Lambiotte, Dansereau, Cross, & Reynolds, 1989;). In the next section we briefly cover some processes involved in reading text displays and we provide an overview of the psychology of reading as presented by Rayner and Pollatsek (1989).

Perceptual Processes in Reading

Reading is the process of extracting information from displays of printed text. This process of information extraction involves making a series of rapid eye movements, known as *saccades*. Between these saccades, the eyes remain fixed on a part of the text for a relatively long time, a process known as *fixations*. Functionally, no useful information is acquired during saccades; when reading text, the visual information available to a viewer consists of a series of discrete snapshots from the fixations (Rayner & Pollatsek, 1989). The fixations following saccades last for approximately 200 to 250 milliseconds, although there is considerable variability in the duration of a fixation. The length of each saccade is influenced by a variety of factors (e.g., the complexity of the text), but it is typically from 7 to 9 letters or spaces. About 10 to 15% of the saccades are regressions in which the eye moves back over text that has already been read. These regressions (or refixations) occur when a word is long or linguistically complex. Readers also skip words that are short or are highly predictable (they

often skip the word *the*). The effective field of view, the *perceptual span*, is affected to some extent by a text's difficulty, font size, and so forth, but it is small under all circumstances. It generally includes approximately 3 or 4 characters to the left of fixation and about 14 to 15 to the right of fixation for English text. The reason for the asymmetry is presumably due to the more informative text lying to the right of fixation; the asymmetry is in the opposite direction for readers of Hebrew, which is read from right to left (Rayner & Pollatsek, 1989). How much information a reader acquires to the right of fixation is related to the difficulty of processing a fixated word; readers acquire less information to the right of fixation when the fixated word is of low frequency. The variability in size of the perceptual span indicates that information outside the high acuity area of the fovea (parafovea) is also used in reading. The length of a fixation on a word is less when the word was previously presented in the parafoveal part of the perceptual span than when it was not.

A key question is: "What information is extracted at the parafoveal area?" Fixation time on a word (e.g., *tune*) is reduced when a visually similar string of letters (e.g., *turc*) has just been presented to parafoveal vision. Presenting a semantically similar word (e.g., *song*) has no effect on subsequent fixation time, although such words produce semantic priming in a lexical decision task. It seems that facilitation is due to abstract (case-independent) letter codes associated with the first few letters of an unidentified parafoveal word. The effect is not due to visual similarity because changing the case of letters from fixation to fixation (e.g., *cHaNgE* presented at fixation and *ChAnGe* presented to the parafovea) has little effect on reading. Consistent with the other context effects that we have discussed, fixation time on a word is reduced if the word is highly predictable on the basis of contextual information. As another example of the influence of meaning on fixation time, the last critical word in clauses and sentences is fixated for a longer duration than is the same word presented in the middle of a clause or sentence. The reason for this result may be due to integrative processes occurring at the ends of clauses and sentences (Rayner & Pollatsek, 1989). Finally, although the same basic mechanisms of eye movement control may occur in reading, scene perception, and visual search, the decision about when to move the eyes may be influenced by different processes for each of these domains (Rayner, 1995). Support for this view was provided by Rayner and Raney (1996), who found that although readers fixate for longer duration on low-frequency words than on high-frequency words, there was no effect of word frequency when subjects searched through texts for a target word. These results indicate that different criteria are used as the basis for executing an eye movement in reading and in visual search.

In this section, we have provided a brief review of the perceptual and cognitive processes involved in reading text-based displays. Research on the factors that determine the process of extracting information from text provided important issues to be considered in developing Web pages. Individual differences in reading have consequences for how effectively information is extracted from text displays. The issues about different decision processes involved in executing eye movements while reading text versus searching and browsing text highlight

the strategic nature of the control of eye movements in visual processing. In the next section, we describe how attention is allocated in processing visual displays and how attention controls visual search processes.

Attention and Visual Search Processes

Many aspects of attention are relevant to the design of Web pages. The WWW is designed to present an abundance of information to users, but only a certain amount of information can be registered in the visual field and selected for further processing at any given time. Because Web users often seek specific information, selective attention is an especially pertinent topic for Web designers. Selective attention refers to the ability to select or process information that is relevant to the present goal while ignoring or inhibiting processing of irrelevant information.

In selecting information for further processing, users must first search the visual field for relevant or target information among distracting or irrelevant sources of information. The visual search process is facilitated by several factors. First, it is easier to detect relevant information when the target information is dissimilar to the irrelevant information. For example, it is easier to detect letters among distractor numbers than to detect letters among letters (Gleitman & Jonides, 1978). If the target information does not have features that are distinctive or unique compared to the distractor information, the time to search the visual display increases as the amount of information in the visual display increases (Duncan & Humphreys, 1989; Treisman, 1986). Detecting a target during a search is easy when the target is a feature singleton (Yantis, 1993); when that is, the target is designated by certain unique or distinctive features (e.g., color, size, orientation). In this case, the unique feature pops-out of the visual display, and detection of the target occurs automatically (Treisman, 1986). Although featural singletons are extremely easy to detect, the designer of a Web page should not overuse them. If there are several feature singletons, the distractor feature singletons can cause users to have difficulty finding the target feature singleton or information not designated by a feature singleton (Theeuwes, 1992). In addition to considerations given to designating target information, the designation of distractor information is also important. It is easier to detect a target among distractors when the distractors are similar to each other than when they are not (Duncan & Humphreys, 1989). For example, the crowding effect mentioned earlier (Anstis, 1974) may partially be due to difficulty identifying letters similar in form. Lowercase letters are more dissimilar in form than are uppercase letters.

Previous knowledge can also influence effective search of a visual display. If a viewer is knowledgeable about where to focus the search in a visual display, he or she can more quickly locate the relevant information. Cuing the location of where relevant information appears, search time decreases but, if the relevant information appears in an unexpected location, search time increases (Posner, Nissen, & Ogden, 1978).

Other previous knowledge that can facilitate visual search is information based on extensive experience or practice in searching a visual display. When a person is initially exposed to a new environment, searching it requires

considerable attention and effort, but as a perceiver becomes experienced with the environment, searching requires less attention and effort. Specifically, if a user consistently searches for the same information and this information is always represented on the visual display in the same way, detecting the information can become automatic (Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977). In designing a Web environment, designers should represent frequently used procedures or search items uniformly and consistently. There is already some consistency in Web page designs: For example, users automatically process underlined text in a different color from text as an informational link that can be followed to gather additional information.

Selective attention can be disrupted when several items on a visual display are designed to capture attention (e.g., color contrasts, movement). In this instance, users may have difficulty focusing on one particular item; in effect, perceivers divide their attention over the visual display. A cost is usually associated with dividing attention: when a perceiver is trying to attend to two items, attention to both items suffers (e.g., Neisser & Becklen, 1975). As mentioned earlier, designers of Web pages should be careful not to overuse mechanisms designed to capture attention so that users have difficulty selecting relevant information and maintaining attention to it. Attention can be divided, however, without any costs to processing more than one item, if processing of these items can occur automatically (Schneider & Fisk, 1982).

CONCLUSION

In our description of visual information processing we have provided a glimpse of the complex processes involved in acquiring information and knowledge about the environment through perception. Although the topics that we have reviewed are essential components of the process of arriving at a percept and acting on the basis of perceptual information, we are aware of other important topics (including motion and color) that we have been unable to cover. Although space has limited our discussion of, these and other topics. We consider the Web to be primarily a text and graphics-based sources of information. Motion and color displays on the Web are typically used to capture and sustain attention, and to facilitate search (particularly when color is used). These are also important aspects of visual information processing, and we have provided some discussion of how displays may be designed to accomplish these purposes. Many principles that we have discussed here can be applied to motion and color.

We have provided a characterization of the environmental information available to perceivers and of their resources for extracting information from the environment. Our aim has been to provide this characterization in such a way as to facilitate the design of information for perceivers. Although perception is robust under most circumstances, we conclude that many factors facilitate perceptual processing. Conversely, some of the same factors that facilitate perception can also impede it. Perhaps the most important conclusion that can be

derived from the present discussion is that perception is a complex process that does not just happen. Unlike many other cognitive skills such as memory, language, and problem solving, perceivers often take their perceptual abilities for granted and sometimes seem to believe that such an easy process cannot be complicated. Designers of information for the Web are in a position to sustain this complacency about perceptual abilities. Designers can control the quality of information available to perceivers so that perceptual processing and user effectiveness are facilitated. We have tried to provide principles of visual information processing that will aid designers in meeting these goals.

REFERENCES

Adams, A. J., Wang, L. S., Wong, L., & Gould, B. (1988). Visual acuity changes with age: Some new perspectives. *American Journal of Optometry and Physiological Optics*, *65*, 403–406.

Anstis, S. M. (1974). A chart demonstrating variations in acuity with retinal position. *Vision Research*, *14*, 589–592.

Besner, D., Smith, M. C., & MacLeod, C. M. (1990). Visual word recognition: A dissociation of lexical and semantic processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 862–869.

Biederman, I. (1983). Scene perception: A failure to find a benefit from prior expectancy or familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*, 411–429.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.

Biederman, I. (1990). Higher-level vision. In D. N. Osherson, S. Kosslyn & J. Hollerbach (Eds.) *An Invitation to Cognitive Science: Visual cognition and action* (pp. 41–72). Cambridge, MA: MIT Press.

Biederman, I., Glass, A. L., & Stacy, E. W., Jr. (1973). Searching for objects in real-world scenes. *Journal of Experimental Psychology*, *97*, 22–27.

Biederman, I., Rabinowitz, J. C., Glass, A. L., & Stacy, E. W., Jr. (1974). On the information extracted from a glance at a scene. *Journal of Experimental Psychology*, *103*, 597–600.

Carr, T. H. & Pollatsek, A. (1985). Recognition of printed words: A look at current models. In D. Besner, T. G. Waller & G. E. MacKinnon (Eds.) *Reading research: Advances in theory and practice* (pp. 1–82). San Diego, CA: Academic Press.

Chen, C., & Rada, R. (1996). Interacting with hypertext: A meta-analysis of experimental studies. *Human-Computer Interaction*, *11*, 125–156.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*, 433–458.

Fisher, K. M., Faletti, J., Patterson, H., Thornton, R., Lipson, J., & Spring, C. (1990). Computer-based concept mapping. *Journal of College Science Teaching*, *19*, 347–352.

Fodor, J. A. (1983). *The modularity of mind.* Cambridge, MA: MIT Press.

Friedrich, F. J., Henik, A., & Tzelgov, J. (1991). Automatic processes in lexical access and spreading activation. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 792–806.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton-Mifflin.

Gillam, B. (1980). Geometrical illusions. *Scientific American*, *242*, 102–111.

Gleitman, H., & Jonides, J. (1978). The effect of set on categorization in visual search. *Perception and Psychophysics*, *24*, 361–368.

Hailman, J. P. (1977). *Optical signals: Animal communication and light.* Bloomington: Indiana University Press.

Hochberg, J. E. (1970). Attention, organization and consciousness. In D. I. Mostofsky (Ed.), *Attention: Contemporary theory and analysis* (pp. 99–124). New York: Appleton-Century-Crofts.

Hoffman, J. E. (1980). Interaction between global and local levels of a form. *Journal of Experimental Psychology: Human Perception and Performance*, *6*, 222–234.

Holley, C. D., & Dansereau, D. F. (1984). Networking: The technique and the empirical evidence. In C. D. Holley & D. F. Dansereau (Eds.) *Spatial learning strategies*. New York: Academic Press.

Julesz, B. (1971). *Foundations of cyclopean perception.* Chicago: University of Chicago Press.

Kimchi, R. (1992). Primacy of wholistic processing and global/local paradigm: A critical review. *Psychological Bulletin*, *112*, 24–38.

Kinchla, R. A., & Wolfe, J. M. (1979). The order of visual processing: "Top-down," "bottom-up," or "middle-out." *Perception and Psychophysics*, *25*, 225–231.

Kosslyn, S. M., & Koenig, O. (1992). *Wet mind, the new cognitive neuroscience.* New York: The Free Press.

Lambiotte, J. G., Dansereau, D. F., Cross, D. R., & Reynolds, S. B. (1989). Multirelational semantic maps. *Educational Psychology Review*, *1*, 331–367.

Land, E. H. (1959). Experiments in color vision. *Scientific American*, *200*, 84–99.

Loftus, G. R., & Mackworth, N. H. (1978). Cognitive determinants of fixation location during picture viewing. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 565–572.

Marr, D. (1982). *Vision: The philosophy and the approach.* New York: W. H. Freeman.

Maunsell, J. H. R., & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience*, *10*, 363–401.

Meyer, D. E., & Schvaneveldt, R. W. (1974). Facilitation in recognizing pairs of words: Evidence for a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227–234.

Nadeau, D. R. (1996). User interface metaphor in virtual reality using VRML. *Behavior Research Methods, Instruments, & Computers*, *28*, 170–173.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, *9*, 353–383.

Navon, D., & Norman, J. (1983). Does global precedence really depend on visual angle? *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 955–965.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. Humphreys (Eds.) *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Lawrence Erlbaum Associates.

Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually significant events. *Cognitive Psychology, 7*, 480–494.

Palmer, S. E. (1975). Visual perception and world knowledge. In D. A. Norman, D. E. Rumelhart, & the L. N. R. Research Group (Eds.), *Explorations in Cognition.* San Francisco: Freeman.

Palmer, S. E., Rosch, E., & Chase, P. (1981). Canonical perspective and the perception of objects. In J. Long & A. Baddeley (Eds.), *Attention and performance*, (Vol. 9, pp. 135–151). Hillsdale, NJ: Lawrence Erlbaum Associates.

Pelli, D. G., Robson, J. G., & Wilkins, A. J. (1988). The design of a new letter chart for measuring contrast sensitivity. *Clinical Vision Science, 2*, 187–199.

Pezdek, K., Whetstone, T., Reynolds, K., Askari, N., & Dougherty, T. (1989). Memory for real-world scenes: The role of consistency with schema expectation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 587–595.

Pinker, S. (1984). Visual cognition: An introduction. *Cognition, 18*, 1–63.

Posner, M. I., Nissen, M. J., & Ogden, W. C. (1978). Attended and unattended processing modes: The role of set for spatial location. In H. L. Pick, Jr. & I. J. Saltzman (Eds.) *Modes of perceiving and processing information*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Potter, M., Moryadas, A., Abrams, I., & Noel, A. (1993). Word perception and misperception in context. *Journal of Experimental Psychology: Learning, Memory and Cognition, 19*, 3–22.

Ramachandran, V. S. (1988). Perceiving shape from shading. *Scientific American, 259*, 76–83.

Rayner, K. (1995). Eye movements and cognitive processes in reading, visual search, and scene perception. In J. M. Findlay, R. W. Kentridge, & R. Walker (Eds.), *Eye movement research: Mechanisms, processes and applications* (pp. 3–22). Amsterdam: Elsevier.

Rayner, K. & Pollatsek, A. (1989). *The Psychology of Reading*. Englewood Cliffs, NJ: Prentice-Hall.

Rayner, K., & Raney, G. E. (1996). Eye movement control in reading and visual search: Effects of word frequency. *Psychonomic Bulletin & Review, 3*, 245–248.

Rock, I. (1983). *The logic of perception*. Cambridge, MA: Bradford Books/MIT Press.

Rock, I., & Palmer, S. (1990). The legacy of Gestalt Psychology. *Scientific American, 263*, 84–90.

Rosch, E. H., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382–439.

Sanders, M. S., & McCormick, E. J. (1993). *Human factors in engineering and design* (7th ed.). New York: McGraw-Hill.

Schneider, W. & Fisk, A. D. (1982). Dual task automatic and control processing, can it be done without cost? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8*, 261–278.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84*, 1–66.

Sekuler, R., & Blake, R. (1994). *Perception* (3rd ed.). New York: McGraw-Hill.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review, 84*, 127–190.

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics, 51*, 599–606.

Todd, J. T., & Reichel, F. D. (1990). The visual perception of smoothly curved surfaces from double projected contour patterns. *Journal of Experimental Psychology: HumanPerception and Performance*, *16*, 665–674.

Treisman, A. (1986). Features and objects in visual processing. *Scientific American*, *255*, 14–125.

Tulving, E., & Gold, C. (1963). Stimulus information and contextual information as determinants of tachistoscopic recognition of words. *Journal of Experimental Psychology*, *66*, 319–327.

Van Essen, D. C. (1985). Functional organization of primate visual cortex. In A. Peters & E. G. Jones (Eds.), *Cerebral Cortex. Volume 3: Visual Cortex* (pp. 259–329). New York: Plenum.

Ward, L. M. (1982). Determinants of attention to local and global features of visual forms. *Journal of Experimental Psychology: Human Perception and Performance*, *8*, 562–581.

Yantis, S. (1993). Stimulus-driven attentional capture. *Current Directions in Psychological Science*, *2*, 156–161.

Yantis, S., & Jonides, J. (1990). Abrupt visual onsets and selective attention: Voluntary versus automatic allocation. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 121–134.

4 Discourse Process and Its Relevance to the Web

Joseph P. Magliano
Michelle C. Schleich
Keith K. Millis
Northern Illinois University

On reading the title of this chapter, many of you may have asked yourselves: "What is discourse processing?" and "What does *that* have to do with the World Wide Web (WWW)?" In this chapter we try to answer these two questions. Although most of the chapter is devoted to answering the first question, we have focused our review of discourse processing on topics that we believe are relevant to Web use and development. We believe that research in discourse processes is useful to Web developers and software engineers wishing to communicate effectively to their audience.

Researchers in discourse processing are interested in how people produce and understand language. It important, though, to distinguish the field of discourse processing from psycholinguistics. Psycholinguistics focuses primarily on how humans understand isolated sentences. For example, a major goal of psycholinguists has been to uncover how readers and listeners derive the syntactic structure of a sentence (see Balota, Flores d'Arcais, & Rayner, 1990). This process involves looking up word meanings in a mental dictionary, breaking the sentence into constituents, assembling the constituents in a tree structure that ultimately indicates who is doing what to whom. Understanding discourse involves much more than parsing sentences in isolation. Rarely does a sentence or utterance occur in isolation; sentences typically occur in a discourse that conveys messages spanning several sentences. Thus, researchers in discourse processing are interested in how connected messages conveyed in conversation and in text are produced, understood, and remembered.

This chapter provides a relatively brief summary of some major distinctions, themes, and problems of interest in the field of discourse processing. This review is by no means complete and those who are interested in learning more about the field are encouraged to read Gernsbacher (1994) and Graesser, Millis, and Zwaan (1996). In particular, we have limited our review to research on text comprehension because we believe that much of this research is useful for to understanding messages conveyed through other media, such as film and multimedia displays.

This review consists of three major sections. The first section describes important characteristics of texts and their corresponding representations in readers' minds. The second section describes current thinking on the nature of text

comprehension. Finally, the third section describes the role of pragmatics in understanding.

CHARACTERISTICS OF TEXT AND TEXT REPRESENTATIONS

Narrative Versus Expository Text

In contrast to narratives, expository texts are typically written to inform readers about a topic. The chapter you are currently reading is an example of an expository text written to inform readers about research in discourse processes. Expository texts include persuasive speech, news, logical arguments, descriptions, and procedural instructions. As with narrative texts, expository texts also have rhetorical structures, but readers may be less familiar with the different rhetorical structures in expository texts than with those in narrative texts. Structure typically depends on the purpose of a text. Procedural texts describe the steps involved in a procedure. In writing such a text, an author may present the information in a list format to emphasize the steps of the procedure. In contrast, a psychology textbook chapter that discusses various models of long-term memory might have the following rhetorical structure: Describe model *X*, and then describe the pros and cons of model *X*.

Readers usually have an easier time comprehending narratives than they do expository texts (Graesser, 1981; Graesser & Riha, 1984). One obvious is that readers are typically more familiar with the material in narratives than they are with that in expository texts. Knowledge about story structure is obtained before school age (Stein & Glenn, 1979), whereas knowledge of the structure of most expository texts is not obtained until much later and might even require formal training. Graesser and Riha (1984) argued that another reason for the advantage of narratives is that the basis for understanding real-life events matches more closely with the contents of narratives than with the content of expositions. Because expository texts do not approximate real life events, they are relatively hard to comprehend.

In our experience with the WWW, we have found that some sites contain stories, others contain expository texts, and still others contain both. Because both types of text are well represented on the WWW, what factors relevant to comprehension are important for Web designers to know when constructing a Web page? One factor is the readers' prior knowledge. When Web designers are preparing a Web page that is expository in nature (i.e., the site is designed to inform readers about a topic), they should not take for granted reader's prior knowledge. Another important factor would be knowing the goal of the reader. A considerate text should be tailored to readers' goals. Hence, a maxim for creating Web text is "know thy reader." Of course, knowing thy reader can be difficult because there are hundreds of ways that people can surf to a particular site. Consequently, a designer might not know readers' prior knowledge or their goals. A brief survey querying visitor goals and knowledge could be used to help refine a Web site.

Text Representations

When readers comprehend text (both narrative and expository), they construct different mental representations reflecting what they have read. One representation, termed the *surface* representation, contains exact wording, font type, and syntactic information. This representation is believed to be short lived; it last only a few seconds after comprehension (Fletcher & Chrysler, 1990). This fact explains why it is often difficult to repeat a sentence verbatim after even a short delay. Another representation, called the *textbase*, is a propositional network that depicts the explicit contents of the text. Propositions are basic ideas that are usually depicted as a predicate connecting two or three arguments. Predicates are usually verbs, and arguments are typically nouns, adjectives, adverbs, or other propositions. For example, the sentence, "This enterprising company wants your business" contains the following propositions: (a) WANTS, COMPANY, BUSINESS; (b) IS, COMPANY, ENTERPRISING; (c) IS, BUSINESS, YOURS. The textbase is incrementally constructed in a network form as the text is read. Many text researchers believe that propositions are connected to each other in the network if they share an argument. This is called the *Principle of Argument Overlap*. According to this principle, proposition (a) would be connected to propositions (b) and (c) because they share the arguments COMPANY and BUSINESS, respectively.

In addition to the surface and textbase representations, readers construct a representation that depicts the situations in the text, called the *situation model*. The situation model contains representations of the characters, events, actions, spatial settings, and objects mentioned in the text and includes inferences readers' world knowledge about the concepts and events described in the text. The situation model can be described as the "mental movie" that people sometimes experience when they read a captivating novel or even a list of procedures such as those detailing how to program a VCR. In contrast to the textbase that represents the propositions inherent in the text, the situation model is similar to what is meant by an understanding of the text.

The extent to which readers construct these representations depends on several factors. First, it depends on the type of text: Because readers have more world knowledge about the topics introduced in narratives than in expositions, they may form stronger and more complete situation models for narratives. Of course, this possibility depends on the overlap of knowledge between reader and text. A scientist reading a text in his or her field of expertise can draw on a potential wealth of knowledge and therefore, can generate a rich situation model for an exposition. In addition, there is evidence that the surface representation is stronger for narratives because readers pay attention to stylistic features of literary text (Zwaan, 1993). Another variable influencing the level of representation is a reader's goal (Zwaan & Oostendorp, 1993). If a person reads to understand a text, than a situation model is stressed, but if a person reads a text to remember it, the textbase is stressed.

The Nature of Understanding

The central focus of discourse-processing research over the past 15 years has been understanding as it is achieved online, that is, during comprehension (see Balota, et al., 1990). Online processes contrast to offline processes that occur sometime after a text is read. Processes that mediate text recall are an example of offline processes. Online and Offline processes are intimately related; long-term memory of a text is mediated by processing that occurs during reading (e.g., Trabasso & van den Broek, 1985). In this section, we discuss some factors that affect both online and off-line text understanding.

Memory Constraints in Understanding. Researchers in discourse comprehension agree that certain psychological constraints and factors influence the construction of the types of representations previously mentioned. One factor is the limited nature of attentional resources allocated to cognitive tasks at a given time. Current models of text comprehension agree that much of the processing necessary to understand a text occurs in a limited working memory (Just & Carpenter, 1992; Kintsch & van Dijk, 1978). Working memory is often characterized as a pool of available resources for the computation and storage of partial products of comprehension. Among other things, working memory provides the resources for retrieving word meanings from the mental lexicon, for abstracting and linking propositions in the textbase, for syntactic parsing, and for generating inferences. Working memory also serves as the storage arena for the outputs of these processes. For example, the sentence The senator the judge wrote died from cancer requires more resources to comprehend than the sentence The judge wrote the senator who died from cancer partly because the former sentence requires comprehenders to store and mentally juggle partial products (e.g., incomplete noun phrases) so that the appropriate noun phrases are linked with the correct verb phrases. Needless to say, people differ in the amount of working memory resources devoted to language processing (Just & Carpenter, 1992). Comprehenders who have efficient language processes (e.g., adult readers) can hold on to more of a current representation in working memory than can comprehenders with less efficient processing abilities (e.g., beginning readers).

The Role of Prior Knowledge in Understanding. Another important factor in comprehension is the knowledge comprehenders have of topics mentioned in a discourse (Bransford & Johnson, 1972; Graesser, 1981; Shank & Abelson, 1977). Psychologists use the term schema to refer to organized packets of knowledge stored in long-term memory (see Mandler, 1984, for a detailed discussion of schema theory). People have schemata about all sorts of things, including objects (e.g., computers, buildings), events (e.g., car accidents, earthquakes), intentional actions (e.g., brushing teeth, going to a job interview), social activities (e.g., dating, birthday parties), and abstract concepts (e.g., morality, truth). In some cases, a person's schema for something is rich and detailed (e.g., going to a restaurant), whereas in other cases, is sparse and incomplete (e.g., nuclear fusion). Comprehension is facilitated to the extent that comprehenders have complete and

rich schemata for the topics mentioned in the text, and those inferred by the comprehender (e.g., Bransford & Johnson, 1972; Shank & Abelson, 1977). Schemata help comprehension by providing a source of top-down processing (DeJong, 1979; Dyer, 1983; Shank & Abelson, 1977). Top-down processing occurs when previous knowledge and experiences affect the processing of incoming information. For example, if a story began, "It was the day for the big birthday party," readers would activate their schema corresponding to birthday party. The information included in this schema aids in interpreting and encoding subsequent information. For example, readers would not be surprised if the story continued "The kids played Pin the Tail on the Donkey" because the concepts of kids and Pin the Tail on the Donkey would be included in the (now activated) birthday party schema. If the birthday party schema was not activated, the interpretation of the sentence would take a longer. Without access to schemata for top-down processing, comprehenders have to resort to bottom-up processing, which involves relying on the properties of a stimulus to arrive at an understanding.

Cohesion and Coherence in Understanding. When readers experience a sequence of text elements (e.g. sentences or clauses), these elements are perceived as being related to each other and part of a whole, rather than a series of isolated elements. The relations that are perceived, however, may be immediate and local and dependent on cohesive elements between adjacent sentences in the text (e.g., argument overlap) or continuities of a situation model (e.g., space, time, and causality). The relations may also be based on global connections between large segments of texts (e.g., distal causal relations and connectivity with thematic goals). Making these local and global connections is central to understanding discourse, especially in view of the limited nature of working memory (Albrecht & O'Brien, 1993; Dahlgren, 1988; Gernsbacher & Givon, 1995; Graesser, Singer, & Trabasso, 1994; Halliday & Hasan, 1976; Kintsch & van Dijk, 1978; Lorch & O Brien, 1995; Singer, Graesser, & Trabasso, 1994; Trabasso & Magliano, 1995).

On what basis do readers establish these local and global relationships between text elements? Before answering this question, first consider the simple narrative shown in Table 4.1. There are several bases by which local or cohesive relationships and global or coherent relationships can be established. Readers can rely on the text (Givon, 1995; Halliday & Hasan, 1976; Kintsch, 1995; Kintsch & van Dijk, 1978) and/or a situational model (Albrecht & O Brien, 1993; Givon, 1995; Magliano, Zwaan, & Graesser, in press; Sanford & Moxey, 1995; van Dijk & Kintsch, 1983; Zwaan, Langston, & Graesser, 1995; Zwaan, Magliano, & Graesser, 1995) to determine whether a sentence fits into local and global contexts.

One way that two successive text elements might be connected is through argument overlap (Fletcher, Chrysler, van den Broek, Deaton, & Bloom, 1995; Kintsch, 1995; Kintsch & van Dijk, 1978). Cohesion is established when the current sentence shares words or phrase with the prior sentence. The degree of argument overlap can vary from none to complete overlap. Consider the example

story "Ivan the Warrior" shown in Table 4.1. Sentences 3 and 4 share three explicit arguments (giant, village, people) in contrast to Sentences 9 and 10 which do not share any. The connection between Sentences 3 and 4 should be perceived as being more cohesive than the connection between Sentences 9 and 10. Some discourse research have regarded argument overlap as the primary basis by which readers establish connectivity between text elements (e.g., Kintsch, 1995; McKoon & Ratcliff, 1992). However, most theories in discourse processing assume that coherent understanding requires more than establishing referential overlap (e.g., Graesser, et al., 1994).

The situation model also provides another basis for establishing local and global connections between text elements. Zwaan and his colleagues have proposed an event indexing model to account for how readers establish situational relationships between narrative story events (Magliano et al., in press; Zwaan, Langston, & Graesser, 1995; Zwaan, Magliano, & Graesser, 1995). This model assumes that the primary function of a situation model is to provide a representation of who is involved in a story, what is happening to them, their goals in context of these events, and what plans they have to achieve their goals, all within narrative time and space. In order to construct this representation, readers monitor the continuity of story events along a number of dimension: time, space, causality, and intentionality.

Events are temporally continuous with the prior context if they occur within the same time frame as the immediate prior context in narrative time. Some story event are discontinuous with the prior context. In the example story in Table 4.1, Sentences 3 and 4 are temporally continuous because they depict successive narrative events (Ivan is being told about the giant), whereas Sentences 12 and 13

Table 4.1
Ivan the Warrior

STORY SENTENCE
1. Ivan was a great warrior.
2. He was the best archer in his village.
3. One day, Ivan heard that a giant had been terrifying people in his village.
4. They said that the giant came to the village at night and hurt people.
5. Ivan was determined to kill the giant.
6. He waited until dark.
7. When the giant came, Ivan shot an arrow at him.
8. Ivan hit him but the arrow could not hurt the giant.
9. The people were disappointed.
10. One day, a famous swordsman came to a nearby village.
11. Ivan decided to learn how to fight with a sword.
12. He went to the swordsman.
13. Ivan studied very hard for several weeks.
14. He became a very skilled swordsman.
15. Ivan got a powerful sword from his teacher.
16. That night, Ivan returned back to his village with his mighty sword.
17. He finally killed the giant with his sword.
18. The people thanked Ivan a hundred times.

are not temporally discontinuous. A story event is spatially continuous with the prior context if it occurs within the same spatial location in narrative space. For example, Sentences 1 through 9 all occur within in Ivan's village, and therefore are related to one another along the spatial dimension. However, Sentence 10 is discontinuous with the prior context along the spatial dimension because it describes an event occurring in a new spatial location (e.g., a nearby village). Story events are connected along the causal dimension if one event provides a reason or explanation for the second event. It is beyond the scope of this chapter to specify how these causal relationships are established (see Trabasso, van den Broek, & Suh, 1989). However, examples from the story may make this dimension more clear. For example, Ivan's goal to kill the giant (Sentence 5) can clearly be explained by the fact that he heard about a giant that had been terrorizing and hurting the people in his village (Sentences 3 and 4). However, there is nothing in the prior context which explains why the swordsman came to the nearby village (Sentence 10). Finally, character actions are continuous along the intentionality dimension if they fit into a pre-established or readily inferable goal plan. Given that character goals and goal-plans are central to narratives, this dimension provides an important basis for establishing global coherence (e.g., Graesser et al. 1994; Trabasso et al., 1989). For example, Ivan's goal to kill the giant provides a basis for understanding his actions throughout the story. There is considerable evidence that monitoring character goals is central to narrative understanding (Albrecht & O'Brien, 1993; Long & Golding, 1993; Long, Golding, & Graesser, 1992; Suh & Trabasso, 1993; Trabasso & Suh, 1993).

It is beyond the scope of the present chapter to provide a detailed discussion of the evidence in favor of the event-indexing model, but we can summarize some general findings. When readers read sentences and are able to achieve cohesion or coherence by integrating sentences into a representation, comprehension is facilitated; when cohesion or coherence fails to be established, comprehension suffers (Albrecht & O'Brien, 1993; Glenberg & Langston, 1992; Graesser, Singer, & Trabasso, 1994; O'Brien & Albrecht, 1992; Singer, Graesser, & Trabasso, 1994; Suh & Trabasso, 1993; Trabasso & Suh, 1993; van den Broek & Lorch, 1993). Sentences consistent with the previous context are processed faster (Albrecht & O'Brien, 1993; Trabasso & Suh, 1993; Zwaan, Magliano, & Graesser, 1995), more integrated in memory (Zwaan, Langston, & Graesser, 1995), remembered better (Fletcher & Bloom, 1988; Goldman & Varnhagen, 1986; Graesser & Clark, 1985; Trabasso & van den Broek, 1985), included more often in summaries (Trabasso & van den Broek, 1985), rated as more important (Trabasso & Sperry, 1985) or necessary and sufficient (Trabasso, van den Broek, & Suh, 1989) than are sentences that are not consistent with the previous context. Furthermore, an entire text is more memorable when its elements are causally cohesive and coherent than when they are not (Trabasso, Secco, & van den Broek, 1984).

Organizational signals provide the final source that readers to rely on when they establish coherence. Organizational signals are writing devices that emphasize the structure of a text (Lorch, 1989). Examples of organizational

signals include outlines, topical overviews and summaries, and headings. Each of these devices can be used to label the major text topics covered in a text. We, for example, have utilized section headings in the present chapter to sift the topics that are covered.

What impact do organizational signals have on online understanding and long-term memory for a text? With respect to online processes, it is reasonable to assume that organizing signals will facilitate online processing by explicitly indicating which information is central to an overall message. Readers can adjust their reading strategies in light of the signals. For example, Lorch and Lorch (1986) found that readers spend more time reading a sentence when it is signaled as containing important information than when it is not. Furthermore, the impact of signals on reading speed is much greater for poor comprehenders than good comprehenders, a finding that suggests that poor comprehenders rely more heavily on organizational signals to guide the allocation of resources during reading than do good comprehenders (Lorch & Lorch, 1986). With respect to long-term memory for a text, organizational signals apparently facilitate the access and recall of its topic structure (Lorch, & Lorch, 1996). Although signaled information is remembered better, there is a cost in that there is a decrease in recall of unsignaled information (Lorch & Lorch, 1986; 1996). Thus, it is important to signal only information that is central to the topic at hand.

If that making connections is important to understanding, then it should be possible to improve comprehension by supplying the connections to readers. For example, Britton and Gulgoz (1991) used a model of text comprehension proposed by Kintsch and van Dijk (1978) in revising a set of expository texts to make them more comprehensible. The Kintsch and van Dijk model provided a basis for determining problematic referring expressions (e.g., unclear pronominal and noun phrase anaphors) and breaks in coherence occurring in the sample of texts used by Britton and Gulgoz (1991). They revised the texts by making the referring expression and local causal relationships explicit. The revisions led to dramatic improvements in comprehension and memory for the texts. One may suspect that these model-based revisions could probably have easily been made through intuition. However, these revisions led to better understanding and memory than did revisions made by writers for *Time* magazine.

When writing a Web page (or any text for that matter), it is important to keep in mind the type of text and the amount of users' knowledge when deciding the extent to which the connections among sentences should be explicit. Although the study by Britton and Gulgoz (1991) suggested that making text connections explicit aids comprehension, it can hurt when overdone. Making explicit *all* possible connections lessens aesthetic appreciation (Roberts & Kreuz, 1993), understanding (Roberts & Kreuz, 1993), and text memory (Myers, Shinjo, & Duffy, 1987), at least for narratives. This finding is not too surprising because if a text makes all connections explicit, it states information that readers readily infer and it becomes boring and tedious to read. This effect may depend on how much a reader knows about the text topic. For readers with low knowledge of a topic, making the connections explicit facilitates comprehension, but for readers with

high knowledge, making the connections explicit can hurt comprehension slightly (McNamara, Kintsch, Songer & Kintsch, in press). Apparently, experts on a topic need to be challenged by a text to some extent to be sufficiently engaged by it.

Inferences and Understanding. When people understand discourse, they often fill in the blanks with inferences drawn from their world knowledge. For example, when reading Sentence 1 from the story in Table 4.1, readers can infer that Ivan is Russian. The basis for this inference presumably is that Ivan is a typical Russian name. Another example of a knowledge-based inference can be generated from Sentence 12 (Ivan went to see the swordsman.): Ivan intends to kill the giant.

The extent to which readers generate knowledge-based inferences during comprehension has been a central issue for discourse-processing researchers (see Graesser & Bower, 1990; Graesser et al., 1994; McKoon & Ratcliff, 1992). Most discourse-processing research has focused on the categories of inferences that are routinely generated online during comprehension. This question is important for writers: If some inferences are not routinely generated by readers, then a writer must make this information explicit in the text (if the writer wishes to pass this information to readers). Researchers in discourse processing and linguistics have proposed several potential taxonomies of inference categories that can be routinely generated during comprehension (e.g., Magliano, Baggett & Graesser, 1996; Singer, 1988). For the present purpose, we discuss a simplified taxonomy of broadly defined inference categories that are functionally distinct from each another. This taxonomy consist of anaphoric inferences, explanations, associations, and predictions.

Readers generate an anaphoric inference when they find the referent to a referring expression (e.g., a pronoun referring to a previous noun). They locate referents to establish cohesive connections between sentences. For example, the story in Table 4.1 contains many pronouns. Readers must find the referents of these pronouns to establish the cohesive relations between the story sentences. Another kind of anaphoric noun reference appears in Sentence 9: Readers must infer that the "people" refer to the same people mentioned in Sentences 3 and 4.

There is considerable evidence that anaphoric inferences are generated during comprehension (Clark & Haviland, 1977; Clark & Sengul, 1979; Corbett & Chang, 1983; Dell, McKoon, & Ratcliff, 1983; Singer, 1980; Singer, Revlin, & Halldorson, 1990). Readers normally activate the referent to a pronoun before they reach the end of the sentence containing the pronoun (Corbett & Chang, 1983) and mentally deactivate other potential referents (Gernsbacher, 1990). Explanations provide reasons why something occurs. They can provide a basis, rationale, motive, idea, condition, occasion, pretext, or why and wherefore of events, states, and actions in a narrative. To understand Sentence 9 of the story in Table 4.1, readers may explain that the people were disappointed because of Ivan's failure to kill the giant who had been terrifying them. These explanatory inferences are based on previous information contained in the text. Explanatory inferences can also be based on general world knowledge. For example, a reader

may explain that the giant is terrifying the people in the village because of the giant's inferred goal to eat the villagers.

Explanatory inferences provide the primary basis for understanding texts as well as events in everyday life (Chi, Bassok, Lewis, Reimann, & Glaser, 1989; Graesser, Bertus, & Magliano, 1995; Graesser et al., 1994; Hilton & Slugowski, 1986; Ranney & Thaggard, 1988; Schank, 1986; Trabasso & Suh, 1993). Explanations may serve to guide comprehension: Readers seem to have a need to know the causes and reasons for events (e.g., Anderson, 1976) and therefore, they attempt to determine *why* something occurs (Graesser et al., 1994, 1995). Furthermore, explanations provide a basis for achieving coherence in understanding (Trabasso & Magliano, 1995; van den Broek, Risden, & Husebye-Hartmann, 1995).

Associations provide information about who does what to whom with what, when, and where. They also provide information about features, properties, relations, and functions of people, objects, and concepts. Associations serve to enrich and fill in the details of a situation model. They offer information concurrent with the information in a focal sentence. For example, when reading that Ivan is a great warrior (Sentence 1), readers may infer that he is strong, wears armor, and fights with a sword. In this example, these associations specify traits of Ivan and the tools that he uses on the basis of readers' general knowledge of what it means to be a warrior.

Considerable research has focused on whether associations are generated during comprehension. Examples of associations that have been studied are instruments necessary to perform an action (e.g., a broom to sweep a floor; cf. Corbett & Dosher, 1978); instantiation of instances for a category (e.g., eating breakfast leads to an inference that someone ate bacon and eggs); the spatial layout of a scene depicted in a story (where a table is located relative to an agent moving through a room; cf. Morrow, Bower, & Greenspan, 1990). Empirical findings, however, suggested that associative inferences are usually not generated online during normal reading (for reviews, see Graesser et al., 1994; Magliano & Graesser, 1991).

Predictions provide the causal consequences of an event. Predictive inferences reflect readers' expectations about possible events, goals, actions, outcomes, and emotions. A prediction may anticipate an event that later occurs in the text and thereby can be substantiated (i.e., confirmed). Although a prediction may be plausible in the context of the focal sentence and the circumstances of the story, it may not be explicitly substantiated by the text. Predictions from a sentence which are later confirmed by the text function to integrate sentences and to create coherence. When reading that the giant has been terrifying and hurting the people in Ivan's village (Sentences 3 and 4), readers can predict that Ivan will want to do something about it. This prediction, of course, is substantiated in Sentence 5.

Much early research suggested that readers do not normally generate predictive inferences during comprehension (Magliano, Baggett, Johnson, & Graesser, 1993; Millis & Graesser, 1994; Potts, Keenan, & Golding, 1988); at best

they generate general predictions, such as "Something bad will happen" (McKoon & Ratcliff, 1986). Several recent studies, however, have shown that predictions are generated when they are causally constrained by the story context (Fincher-Kiefer, 1993; Murray, Klin, & Myers, 1993), when the eliciting information is foregrounded in the story (Keefe & McDaniel, 1993; Magliano, Dijkstra, & Zwaan, 1996; Whitney, Ritchie, & Crane, 1992), or when the prediction resolves a coherence break (Murray et al., 1993).

Of the inference categories discussed in this brief review, the preponderance of evidence support they idea that only anaphoric referents and explanations are routinely generated online, whereas associations and predictions are generated only under certain circumstances. Most research supported the theoretical position that understanding is largely explanation based (e.g., Graesser, et al., 1994).

The Role of Pragmatics in Understanding

Our final topic in discourse processes involves pragmatics in understanding. Pragmatics refers to knowledge of the social rules that underlie spoken and written communication. Researchers who study pragmatics are interested in the ways speakers/listeners and writers/readers rely on these rules to communicate effectively. When they follow the rules, comprehension is facilitated. When they are violated them, comprehension suffers (Clark, 1993; Grice, 1975). In this section, we discuss some well documented ground rules of communication.

One important component to effective communication is the monitoring of *common ground*, the shared knowledge between speakers (or between writers and readers). To communicate effectively, a speaker/writers must assess listener/readers' knowledge of specific words, ideas, and experiences throughout the course of a conversation or discourse. When there is little common ground, the speaker (or writer) must be more explicit than when there is more shared knowledge.

Another principle of effective communicators is the given-new strategy (Clark & Haviland, 1977; Haviland & Clark, 1974). When people communicate through conversation or written text, they must effectively manage the presentation of information that has already been discussed (i.e., given) and information that has not been previously discussed (i.e., new). The given-new strategy is a simple communication heuristic that good speakers and writers follow to assure that new or given information is clearly marked. Given information is typically included in the subject noun phrase of a sentence or in the first clause of a multiclause sentence, and new information is included in the object clause or the additional clauses (Clark & Haviland, 1977; Haviland & Clark, 1974).

Another set of simple rules of good communication was specified by Grice (1975), who proposed a set of *maxims* that facilitate cooperative communication when they are adhered to by speech participants. Under normal circumstances, participants assume that both parties are adhering to these maxims, but if a maxim is violated by a participant, the cooperative nature of the conversation is

sacrificed. These maxims have been typically addressed in the context of conversations, but they are relevant for writers and Web designers because readers and writers are essentially engaged in dialogues. The maxims introduced by Grice (1975) are:

> *Quantity*: Make your contribution as informative as required, but not any less or more so.
> *Quality*: Make your contribution true — do not lie or say things for which you lack evidence.
> *Relation*: Make your contribution relevant to the discourse.
> *Manner*: Avoid ambiguity and obscurity of expression.

We found violations of these maxims on the WWW. One paragraph in a Web page explained to readers in excruciating detail how to send E-mail when only one sentence was needed. Another included only a simple instruction ("click here") that should have had some corresponding text to explain why readers should click there. For most readers, these examples violated the maxim of quantity. Another Web page promised to make readers millionaires without including the details and the fine print; this strategy violated the maxim of quality. Yet another Web page presented a hodge podge of topics and links without really specifying the page's theme. Another included embedded hypertext links that sent us to Web cites that were only tangentially related to the main topic. These pages violated the maxim of relation. Last, a Web page discussing radium was scattered with scientific jargon, but the writer had made not effort to translate its jargon. For us, this page violated the maxim of manner.

A violation for one reader may not be a violation for another: Readers' knowledge of the common ground between reader and writer affect comprehensibility. A scientist may easily read radium text, whereas others might not. Thus, writers should make clear to readers where to learn about the topics to be discussed on a page before they read further. This information could consist of links or embedded hypertext. As already mentioned, it is important for readers to know the content of an embedded link before they select it.

CONCLUSION

We started this chapter by posing two questions: "What is discourse processing?" and "What does *that* have to do with the WWW?" The present review provides a reasonable answer to the first question. Discourse-processing researchers are interested in how understanding is achieved when people comprehend connected discourse. We discussed the major factors that play a role in achieving understanding. The second question is more difficult to answer than the first, but we think that the research discussed here is relevant to designing Web pages and that Web designers can benefit from knowledge of discourse-processing research. The usefulness of discourse-processing research to Web design has yet to be

determined, however. We also think that it is time for discourse- processing researchers to study how understanding is achieved in alternative text formats such as the WWW. Although we have argued that the nature of understanding is similar across contexts (e.g., reading a novel, watching a movie, or surfing the WWW), this claim must be empirically validated.

REFERENCES

Albrecht, J. E., & O'Brien, E. J. (1993). Updating a mental model: Maintaining both local and global coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 1061-1070.

Anderson, J. R. (1976). *Language, memory, and thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Balota, D. A., Flores d'Arcais, G. B., & Rayner, K. (Eds.) (1990). *Comprehension processes in reading*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Bransford, J., & Johnson, M. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning and Verbal Behavior, 11*, 717-726.

Britton, B. K., & Gulgoz, S. (1991). Using Kintsch's computational model to improve instructional text: Effects of repairing inference calls on recall and cognitive structures. *Journal of Educational Psychology, 83*, 329-345.

Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, R., & Glaser, R. (1989). Self-explanation: How students study and use examples in learning to solve problems. *Cognitive Science, 13*, 145-182.

Clark, H. H. (1993). *Arenas of language use*. Chicago: University of Chicago Press.

Clark, H. H., & Haviland, S. E. (1977). Comprehension and the given-new contract. In R. O. Freedle (Ed.) *Discourse production and comprehension* (pp. 1-40). Norwood, NJ: Ablex.

Clark, H. H. & Sengul, C. J. (1979). In search of referents for nouns and pronouns. *Memory and Cognition, 7*, 35-41.

Corbett, A. T., & Chang, F. R. (1983). Pronoun disambiguation: Accessing potential antecedents. *Memory and Cognition, 11*, 283-294.

Corbett, A. T., & Dosher, B. A. (1978). Instrument inferences in sentence encoding. *Journal of Verbal Learning and Verbal Behavior, 17*, 479-491.

Dahlgren, K. (1988). *Naive semantics for natural language understanding*. Boston: Kluwer Academic.

DeJong, G. (1979). Prediction and substantiation: A new approach to natural language processing. *Cognitive Science, 3*, 251-273.

Dell, G. S., McKoon, G., & Ratcliff, R. (1983). The activation of antecedent information during the processing of anaphoric reference in reading. *Journal of Verbal Learning and Verbal Behavior, 22*, 121-132.

Dyer, M. G. (1983). *In-depth understanding: A computer model of integrated processing for narrative comprehension.*. Cambridge, MA: MIT Press.

Fincher-Kiefer, R. (1993). The role of predictive inferences in situation model construction. *Discourse Processes, 16*, 99-124.

Fletcher, C. R., & Bloom, C. P. (1988). Causal reasoning in the comprehension of simple narrative texts. *Journal of Memory and Language, 27*, 235-244.

Fletcher, C. R., & Chrysler, S. T. (1990). Surface forms, textbases, and situation models: Recognition memory for three types of textual informaton. *Discourse Processes, 13*, 175-190.

Gernsbacher, M. A. (1990). *Language comprehension as structure building*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gernsbacher, M. A. (Ed.). (1994). *Handbook of psycholinguistics*. New York: Academic Press.

Gernsbacher, M. A., & Givon, T. (Eds.). (1995). *Coherence in spontaneous text*. Philadelphia: Benjamins.

Givon, T. (1995). Coherence in text vs. coherence in mind. In M. A. Gernsbacher and T. Givon (Eds.), *Coherence in spontaneous text* (pp. 189-214). Philadelphia: Benjamins.

Glenberg, A. M., & Langston, W. E. (1992). Comprehension of illustrated text: Pictures help to build mental models. *Journal of Memory and Language, 31*, 129-151.

Goldman, S. R., & Varnhagen, C. K. (1986). Memory for embedded and sequential story structures. *Journal of Memory and Language, 25*, 401-418.

Graesser, A. C. (1981). *Prose comprehension beyond the word*. New York: Springer-Verlag.

Graesser A. C., Bertus, E. L., & Magliano, J. P. (1995). Inference generation during the comprehension of narrative text. In E. J. O'Brien & R. F. Lorch (Eds.), *Sources of coherence in text comprehension* (pp. 295-319). Hillsdale, NJ: Lawrence Erlbaum Associates.

Graesser, A. C., & Bower, G. H. (Eds.). (1990). *The psychology of learning and motivation: Inferences and text comprehension*. San Diego, CA: Academic Press.

Graesser, A. C., & Clark, L. F. (1985). *Structures and procedures of implicit knowledge*. Norwood, NJ: Ablex.

Graesser, A.C., Millis, K. K., & Zwaan, R. A. (1996). Discourse comprehension. *Annual Review of Psychology, 48*, 163-189.

Graesser, A. C., & Riha, J. R. (1984). An application of multiple regression techniques to sentence reading times. In D. Kieras & M. Just (Eds.), *New methods in comprehension research*. (pp. 183-218). Hillsdale, NJ: Lawrence Erlbaum Associates.

Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review, 101*, 371-395.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.) *Syntax and semantics: Speech acts* (Vol. 3, pp. 41-58). San Diego, CA: Academic Press.

Halliday, M. A. K., & Hasan, R. (1976). *Cohesion in English*. London: Longman.

Haviland, S. E., & Clark, H. H. (1974). What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior, 13*, 512-521.

Hilton, D. J., & Slugowski, B. R. (1986). Knowledge-based causal attribution. *Psychological Review, 93*, 75-88.

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review, 99*, 122-149.

Keefe, D. E., & McDaniel, M. A. (1993). The time course and durability of predictive inferences. *Journal of Memory and Language, 32*, 446-463.

Kintsch, W. (1995). How readers construct situation models for stories: The role of syntactic cues and causal inferences. In M. A. Gernsbacher and T. Givon (Eds.), *Coherence in spontaneous text* (pp. 189-214). Philadelphia: Benjamins.

Kintsch, W., & Van Dijk, T. A. (1978). Toward a model of text comprehension and production. *Psychological Review*, *85*, 363-394.

Long, D. L., & Golding, J. M. (1993). Superordinate goal inferences: Are they automatically generated during comprehension. *Discourse Processes*, *16*, 55-73.

Long, D. L., Golding, J. M., & Graesser, A. C. (1993). A test of online status of goal-related inferences. *Journal of Memory and Language*, *31*, 634-647.

Lorch, R. F. (1989). Text-signaling devices and their effects on reading and memory processes. *Educational Psychology Review*, *1*, 209-234.

Lorch, R. F., & Lorch, E. P. (1986). Online processing of summary and importance signals in reading. *Discourse Processes*, *9*, 489-496.

Lorch, R. F., & Lorch, E. P. (1996). Effects of organizational signals on free recall of expository text. *Journal of Educational Psychology*, *88*, 38-48.

Lorch, R. F., & O'Brien, E. J. (Eds.). (1995). *Sources of coherence in reading*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Magliano, J. P., Baggett W. B., Johnson, B. K., & Graesser, A. C. (1993). The time course in which causal antecedent and causal consequence inferences are generated. *Discourse Processes*, *16*, 35-53.

Magliano, J. P., Dijkstra, K., & Zwaan, R. A. (1996). Generating predictive inferences while viewing a movie. *Discourse Processes.*, *22*, 199-224.

Magliano, J. P., & Graesser, A. C. (1991). A three-pronged method for studying inference generation in literary text. *Poetics*, *20*, 193-232.

Magliano, J. P., Baggett W. B., & Graesser, A. C. (1996). A taxonomy of inference categories that may be generated during the comprehension of literary texts. In R.J. Kreuz & M.S. MacNealy (Eds.), *Empirical approaches to literature and aesthetics* (pp 201- 220). Norwood, NJ: Ablex.

Magliano, J. P., Zwaan, R. A., & Graesser, A. C. (in press). The role of situational continuity in narrative understanding. In S. R., Goldman & H. van Oostendorp (Eds.), *The construction of mental representations during reading*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Mandler, J. M. (1984). *Stories, scripts, and scenes: Aspects of schema theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 82-91.

McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review*, *99*, 440-466.

McNamara, D. S., Kintsch, E., Songer, N. B., & Kintsch, W. (in press). Text coherence, background knowlege and levels of understanding in learning from text. *Cognitive Intruction*.

Millis, K. K., & Graesser, A. C. (1994). The time-course of constructing knowledge-based inferences for scientific texts. *Journal of Memory and Language*, *33*, 583-599.

Morrow, D. G., Bower, G. H., & Greenspan S. L. (1990). Situation-based inferences during narrative comprehension. In A. C. Graesser and G. H. Bower (Eds.), *Inferences and text comprehension* (pp. 123-135). New York: Academic Press.

Murray, J. D., Klin, C. M., & Myers, J. L. (1993). Forward inferences in narrative text. *Journal of Memory and Language*, *32*, 464-473.

Myers, J. L., Shinjo, M., & Duffy, S. A. (1987). Degree of causal relatedness and memory. *Journal of Memory and Language, 26*, 453-465.

O'Brien, E. J., & Albrecht, J. E. (1992). Comprehension strategies in the development of a mental model. *Journal of Experimental Psychology: Leaning, Memory, and Cognition, 18*, 1061-1070.

Potts, G. R., Keenan, J. M., & Golding, J. M. (1988). Assessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory and Language, 27*, 399-415.

Ranney, M., & Thaggard, P. (1988). Explanatory coherence and belief revision in naive physics. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (pp. 426-432). Hillsdale, NJ: Lawrence Erlbaum Associates.

Roberts, R. M., & Kreuz, R. J. (1993). Nonstandard discourse and its coherence. *Discourse Processes, 16*, 777-784.

Rumelhart, D. E. (1975). Notes on a schema for stories. In D. C. Bobrow and A. M. Collins (Eds.) *Basic processes in reading: Perception and comprehension.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Sanford, A. J., & Moxey, L. M. (1995). Aspects of coherence in written language: A psychological perspective. In M. A. Gernsbacher and T. Givon (Eds.) *Coherence in spontaneous text* (pp. 189-214). Philadelphia: Benjamins.

Schank, R. C. (1986). *Explanation patterns: Understanding mechanically and creatively.* Hillsdale, NJ: Lawrence Lawrence Erlbaum Associates.

Shank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding.* Hillsdale, NJ.: Lawrence Erlbaum Associates.

Singer, M. (1980). The role of case-filling inferences in the coherence of brief passages. *Discourse Processes, 3*, 185-201.

Singer, M. (1988). Inferences in reading comprehension. In M. Daneman, G. E. Mackinnon, & T. G. Waller (Eds.), *Reading reasearch: Advances in theory and proactice*, (Vol. 6, pp. 177-219). San Diego: Academic Press.

Singer, M., Graesser, A. C., & Trabasso, T. (1994). Minimal or global inference during reading. *Journal of Memory and Language, 33*, 421-441.

Singer, M., Revlin, R., & Halldorson, M. (1990). Bridging-inferences and enthymeme. In A. C. Graesser & G. Bower (Eds.), *Inferences and text comprehension.* New York: Academic Press.

Stein, N. L., & Glenn, C. G. (1979). An analysis of story comprehension in elementary school children. In R. O. Freedle (Ed.), *New directions in discourse processing: Advances in discourse processes* (pp. 53-120). Norwood, NJ: Ablex.

Suh, S., & Trabasso, T. (1993). Inferences during reading: Converging evidence from discourse analysis, talk-aloud protocols, and recognition priming. *Journal of Memory and Language, 32*, 279-301.

Trabasso, T., & Magliano, J. P. (1995). Understanding emotional understanding. In N. Frijda (Ed.) *Proceedings of the International Society of Research on Emotions* (pp. 78-82). Hillsdale, NJ: Lawrence Erlbaum Associates.

Trabasso, T., Secco, S., & van den Broek, P. (1984). Causal cohesion and story coherence. In H. Mandl, N. L. Stein, & T. Trabasso, (Eds.), *Learning, and comprehension of text* (pp. 83-111). Hillsdale, NJ: Lawrence Erlbaum Associates.

Trabasso, T., & Sperry, L. (1985). Causal relatedness and importance of story events. *Journal of Memory and Language, 24*, 595-611.

Trabasso, T., & Suh, S. (1993). Understanding text: Achieving explanatory coherence through online inferences and mental operations in working memory. *Discourse Processes*, *16*, 3-34.

Trabasso, T., & van den Broek, P. (1985). Causal thinking and the representation of narrative events. *Journal of Memory and Language*, *24*, 612-630.

Trabasso, T., van den Broek, P., & Suh, S. (1989). Logical necessity and transitivity of causal relations in the representation of stories. *Discourse Processes*, *12*, 1-25.

van den Broek, P., & Lorch, R. F. (1993). Network representations of causal relations in memory for narrative texts: Evidence from primed recognition. *Discourse Processes*, *16*, 75-98.

van den Broek, P. Risden, K., & Husebye-Hartmann, E. (1995). The role of reader's standards of coherence in the generation of inferences during reading. In R. F. Lorch & E. J. O'Brien (Eds.) *Sources of coherence in reading* (pp. 353-374). Hillsdale, NJ: Lawrence Erlbaum Associates.

van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. New York: Academic Press.

Whitney, P., Ritchie, B. G., & Crane, R. S. (1992). The effect of foregrounding on readers' use of predictive inferences. *Memory and Cognition*, *20*, 424-432.

Zwaan R. A. (1993). *Aspects of literary comprehension*. Amsterdam: Benjamins.

Zwaan, R. A., Langston, M. C., & Graesser, A. C. (1995). The construction of situation models in narrative comprehension: An event-indexing model. *Psychological Science*, *6*, 292-297.

Zwaan, R. A., Magliano, J. P., & Graesser, A. C. (1995). Dimensions of situation model construction in narrative comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 386-397.

Zwaan, R. A. & Oostendorp, H. (1993). Do readers construct spatial representations in naturalistic story comprehension? *Discourse Processes*, *16*, pp. 125-143.

5 Human Navigation

Leslie A. Whitaker
University of Dayton

Navigation is a term that describes activities ranging from the first tentative exploration by an infant to the sophisticated calculations and planning which successfully placed a man on the moon. Navigation in its narrow sense means to move through space; in its broader sense, navigation also includes virtual movement through cognitive space made up of data and the knowledge emerging from those data. It is in this broader sense that navigation through cyberspace needs to be considered by Web site designers. Does the Web site design facilitate easy navigation within itself and among related sites? All navigation requires cognitive skills that allow the navigator to accomplish his or her goals. The purpose of this chapter is to describe those cognitive skills and their relationship to navigating the World Wide Web.

The first issue to be addressed is a person's purpose in using the WWW. Two major categories of purpose are to achieve an end goal (e.g., What is the name of the company which produces replacement needles for my 25 year old phonograph?) and to explore for the sake of exploration (e.g., I wonder what will happen if I click on that icon). Some of the navigational skills (and hence design considerations) are the same for smoothly achieving both of the purposes, but some are unique.

NAVIGATION TOWARD A GOAL

When people navigate toward goals, they generally use one of three methods for doing so. They navigate by landmarks, route finding, or survey knowledge (Wickens, 1992). Each of these methods is used under different circumstances.

Landmarks. This type of navigation employs salient landmarks at intermittent points along a route of travel. When a person navigates by landmarks, it is critical that distinctive markers be used in the display to aid both in finding the route to the goal site as well as in determining where the navigator is currently located along the path to the site. Unique visual cues are the critical element of this navigational method.

Route Knowledge. This navigational method allows the user to string together a series of landmarks that determine the route to be followed. Although unique visual cues are an important element of this method, in addition, it is necessary for the navigator to know the steps to be taken to move between the landmarks. An important aspect of this method is that the terrain and movement through it is seen only from the user's perspective at any moment in time. That is, route knowledge only works if you are on a known route. It is not useful for recovering the route after losing your way. Furthermore, if a particular route is inefficient for reaching a site, the route knowledge is not sufficient to allow the navigator either to see the inefficiency or to find an alternative route that may achieve the same final goal site more efficiently.

Survey Knowledge. Navigators able to employ this method have sufficient knowledge to form a cognitive map of the navigational space (Tolman, 1948). A *cognitive map* is an internal representation analogous to a physical map of the space. Although this cognitive map is not perfectly veridical, it does provide a world-centered frame of reference on which the navigator can base decisions and generate navigational choices. One primary advantage of this world-centered framework is that the navigator can move along routes not previously traveled, recover from being lost, and take shortcuts toward goal sites. It is known that survey knowledge is gained by extensive experience with a particular environment. The evidence is more equivocal about whether there are also innate traits in some people that allow them to use survey knowledge and whether experience with navigation in general (as opposed to that in a specific environment) encourages the formation of survey knowledge about a specific environment.

An additional method of navigation that is available to users of the web but not yet available to navigators in the physical world is the use of search engines such as Yahoo. These search engines bring the power of machine intelligence and memory capability to bear on the problems of moving through the vast information set provided by the Web. However, none of these engines is all powerful and many are very restricted in their type of search strategies or their topics of search. A colleague has described a search engine as a native guide that is very familiar with one particular region of the world. This guide's usefulness is further restricted because he is intimately familiar with only one aspect of that world. A good guide may lead a person on a successful fishing holiday in northern Michigan; however, this same guide would be useless in leading a dolphin watching expedition in Florida. The analogy is that a particular search engine may be able to search only particular databases by using key words; however, its ability to search a relational database would be nonexistent.

With respect to the design of Web sites and links, there is an important message learned from an understanding of individual differences in the use of navigational methods. The lesson learned is just because the navigation of related Web sites seems to be apparent to the designer, an experienced Web surfer, does not mean that it will be apparent or even compatible with that of the intended users of the site(s).

Navigation as Exploration

Alternatively, some navigation is undertaken for the value of the journey itself. John Steinbeck described this type of navigation as educational travel; that is, the *process* of moving through space is its own goal. Although only a few navigators are able to traverse the physical world as dilettantes, many users surf the net with the primary purpose of enjoying the process of exploration. Therefore, the design of Web sites and their links should take this exploration function into consideration. Design aspects that permit the user to orient in both space and time allow the explorer to maintain an important feature known as situation awareness. Dominguez (1994) defined situation awareness as continuous extraction of environmental information, integration of this information with previous knowledge to form a coherent mental picture, and the use of that picture in directing further perception and anticipating future events. This concept is one of the more promising ones in the field of applied cognitive psychology and it encompasses a critical skill that a user must maintain while exploring the Web.

Investigations that have examined the maintenance of situation awareness among aircraft pilots have concentrated on the experience and knowledge of the pilots. It may be beneficial for a Web site designer to consider the importance of navigators' maintaining situation awareness and attempt to determine what design features facilitate this.

Exploration may lead a person via an unintended path to a site of interest. When this occurs, two design features help subsequent navigation to the same or a related Web site. These two features, already employed by successful Web browsers, are a history of sites visited and a bookmark.

A history of sites is usually contained in a pull down menu and provides a "trail of breadcrumbs" to mark the path of Web pages through which the user has recently moved. This information serves two navigational functions: it tells the user the landmarks along the path to reach the present location — this function is known as route finding for navigators in the physical world — and it allows the user to click several pages back without having to retrace each page of the route. This function serves as an electronic shortcut which does not have an analog in the physical world.

Bookmarks are placeholders created by the user for future use. This is a function provided by some Web browsers and is not a feature that must be specifically designed into a Web site. Bookmarks are retained after an internet session is finished. In this way, the user can return to a useful site without following the entire route to get to that site; this function serves as an electronic teleport to the location of interest.

Navigating in Structured Environments

The majority of research into human navigation is based on navigation in structured environments (e.g., towns, shopping malls, hospital corridors, etc.). These environments have two general features that influence a person's ability to

navigate through them: unique features that can serve as landmarks and a standard structure for each such environment.

Thorndyke (1980) discussed the importance of unique landmarks for navigators who are new to a city. As described in the previous sections, navigation in a new environment is primarily conducted through the use of landmarks. This is a critical concept for the designer of a Web site. The use of a unique icon to mark a function is employed by the user to navigate through the functions of the Web (page or links). The development and acceptance of such unique and standardized icons is important for improving the navigability of the Web, particularly by new or occasional users. Considering such functional place holders to be navigational landmarks aids the designer in understanding their application by the user. A further consideration for designers is the placement of these icons. If an arrow is adopted as the "turn the page" icon, then it should be placed in a consistent position on the Web page. Again, an analogy to city navigation may be helpful. The design of traffic control devices requires not only that all cities accept that a red signals stop, but also that the red signal be located in approximately the same place in relation to the driver's visual scanning pattern. Even the change from placement on a vertical pole to placement on a horizontal wire above the street (e.g., in Washington, DC) has caused some drivers to complain that they did not see the light. The MIS design military standard (Human Engineering Guidelines, 1985) provides an example of placing functions in a consistent location. Probably starting with a DOS interface, the standard recommends that all interfaces provide a location in the lower left corner of the screen for the presentation of working status of the machine. In this way, the user knows to look in a single place for status information.

Again, considering Thorndyke's developmental hierarchy, navigation in a built environment becomes much more efficient as the navigator learns the general layout of the city and can move from a path linking a series of landmarks to a plan that encompasses several different ways to achieve the same results. Thus, a more experienced navigator is able to move about the Web more efficiently when the Web contains a consistent functionality in the layout of its various parts. It is worthwhile to note that the experienced Web user and the Web site designer are much more similar in their navigational styles than are the novice user and the designer. Thus, it is far easier and more natural for the designer to meet the navigational needs of this user than the needs of the inexperienced user.

Navigating in Unstructured Environments

Less research has been reported that describes navigation through natural environments rather than through structured environments. However, some studies of orienteering and cross-country exploration have examined the strategies used by navigators in off-road environments, as well as the cues provided by those environments. To the extent that the Web allows the user to employ similar strategies and cues in navigating the Web, results reported by these studies are of interest.

Whitaker (1996) cautions that navigation in structured environments is qualitatively different from navigation in natural ones. These differences include both the cognitive processes required and the environmental cues used. This makes the issue of Web site design a particularly tricky one because at one instance the user may think and behave as though the Web were a structured environment and in another instance as though it were a quasinatural one, grown like topsy with no particular overall pattern or structure beyond that imposed by space and available information.

There are several factors that may help the designer anticipate which frame of reference a user may employ in any given site. Structured environments like cities differ from those of naturally constructed environments like off-road areas in the following ways:

- Structured environments make extensive use of straight lines and square corners. The presence of these physical markers are sufficient to designate the environment as a structured one and to shift the navigator's frame of reference to one appropriate for such environments.
- Structured environments make extensive use of signing and linguistic cues. Streets are signed; maps use names of features and these names are actually found in the environment. In this way, a navigator is able to make use of signs, even when their meaning is unknown. For example, navigation in Japan often requires the use of signing written in kanji. To navigate successfully, a person does not need to know the meaning of the symbols, but can merely match the kanji on a sign with the same symbols written on a map. Similarly, navigation in Moscow can be accomplished without knowledge of Russian place names, but by matching a place name written on paper with the displayed names on the street signs.
- Unstructured environments allow the navigator to move more slowly and to pace the speed of that movement. In contrast, structured environments, particularly roadways or public transportation often dictate the speed at which the user must navigate them. That is, self-pacing may be lost when navigating through a structured environment.

Keeping these distinguishing characteristics in mind helps the designer to know whether his user is more likely to navigate through the Web site like a trip to a point of interest located in a new city or like an expedition in the mountains looking for an unmarked gold mine.

Whitaker and Cuqlock-Knopp (1992) reported a taxonomy that described both the strategies and the cues used by navigators in natural (unstructured) environments. This taxonomy was developed through critical incident interviews with experienced off-road navigators. The cognitive strategies in the taxonomy are probably of greater interest to designers than are the types of cues because the strategies describe the cognitive model by which the navigator moves through an unstructured environment, much as Thorndyke's models describe the cognitive models guiding navigation for structured environments. The environmental cues

used by navigators in the physical unstructured world will first be described briefly to provide a content on which the navigational strategies work.

The four types of visual cues used by off-road navigators are manmade, land contours, water features, and vegetation. In unstructured environments, manmade cues are rare; however, they are so distinctive that these experienced navigators reported that they were very valuable. Land contours had value because they are permanent, often correctly mapped, indicative of the difficulty of the terrain, and can be seen from a distance. Water features such as streams and lakes were important because of their relative permanence and their impediment to travel. Finally, vegetation was the most transient of the features used; vegetation was often incorrectly mapped because of the time between the mapping and the moment of navigation. However, vegetation provided information about the ease of movement and was therefore important.

The strategies employed by off-road navigators can be best understood in the context of navigating the Web as problem-solving strategies that apply to a variety of unstructured environments, whether physical or virtual. Four such strategies were included in the taxonomy of these navigators: prediction, recovery, use of catching features, and aiming off.

1. *Prediction* means the ability to anticipate what location will appear next. For example, "There is a stream crossing around the next bend of this trail." The problem-solving strategy uses such predictions to determine whether they have strayed from their intended route. That is, the stream-crossing prediction above is used as follows: "If I'm right about my location, then I should find a stream crossing around this next bend."
2. *Recovery* describes the problem-solving process that must be followed to allow a navigator to recover from being lost. Experienced users reported that when they were lost they did not continue to blindly thrash across the countryside, hoping to see some feature that would orient them. Some would physically retrace their steps to their last known positions; alternatively, others would mentally retrace their actions to develop a hypothesis of where they currently were. A final recovery strategy was to engage in predication in the hopes of discovering a position from which they could determine their actual location.
3. *Catching features* are features that guard the navigator from going too far beyond their goal location. In structured environments, a catching feature is "If you reach the high school, you've gone two blocks beyond my street." In an unstructured environment, catching features are available too; however, they must often take a different form because the route of the navigator is less constrained in an unstructured environment. In the physical world, these features were often ones that provided a long barrier to further travel (e.g., a cliff or a stream).
4. *Aiming off* is the strategy of deliberately planning a route that does not bring the navigator to the exact goal location. This strategy is used in combination with a long natural feature that can be used to guide the navigator to the

desired location. An example using a stream ford (i.e., low water river crossing) may help to clarify this point. Suppose that a navigator knows that a ford exists due north of her present location. She can set her compass on north and hope to accurately follow this heading. If she errs and misses the ford, she will not know whether to turn east or west to follow the stream to the ford. Alternatively, she can aim off by deliberately following a heading of 15 degrees east of north. This route brings her to the stream clearly east of the ford. A left turn and a short jog along the stream brings her to the ford.

NAVIGATING ON THE WEB

The Web provides both a mental space and a virtual physical environment through which a human can move either to achieve a goal or to enjoy the pleasure of "educational travel." The users of the Web page may be very experienced or they may have only rudimentary knowledge of computers. Creating a system that will meet these diverse needs is a tall order for any designer. However, there are existing web sites which use navigational concepts wisely and provide examples for the designer. The Visible Human Project at 'http://www.nlm.nih.gov/research/visible/visible_human.html' describes a physiological model of human anatomy. This site makes good use of a consistent look and feel; standard location of control icons, key words to aid the initial search, and active text to allow movement among the pages of the site.

Navigation Design Principles

The design choices that the Web master makes in constructing a Web site and its links influence the ease and success of its use. As a starting point for the designer, this final section provides a summary of the material described in this chapter. The summary is written in the form of a list of navigation design principles. These principles have not yet been tested in the building of Web sites. In fact, the purpose of this book (not this chapter) is to develop an inclusive set of guidelines. Therefore, the following "principles" should be taken as the result of a thought experiment applying research results from navigation in the physical world to navigation in cyberspace. It awaits your confirmation or refutation as the actual designers of Web pages.

Summary of Chapter as Design Principles

Principle 1. Determine why a user will visit your web site. What are the expected links for arriving and leaving the site? What is the user's purpose in viewing the site?

Principle 2. Learn the proficiency level of the users. Are they following a string of landmarks to move among options (site links as well as within a site)? Do they have alternative routes for reaching your site if one of their links is nonfunctional or altered? Are they experienced users with good survey knowledge of the Web?

Principle 3. Know the search engines used by most of your users. Do the users generally have a search engine permitting the power of machine intelligence and memory to search for your site? What are the limitations of those search engines? Will a user find your site if the search strategy of the software is altered? How can you make your site more accessible by adding automatic links that make use of these engines?

Principle 4. Design to allow the user to maintain situation awareness. Are the users able to maintain situation awareness while moving to and through your web site? Do they know how deep they are in the site? Can they relate the information they are currently seeing with where they came from and what their goal site is?

Principle 5. Make constructive use of landmarks to aid revisiting your site. Does your site make good use of existing audit trails (history of sites visited) and bookmark by using distinctive titles that add recall? Remember that each page can be selected as a bookmark. If users happen on your site during the course of exploration, can they find you again on a later visit? Do your pages have distinctive titles (e.g., *Joe's Zoo Page* instead of *Page 6*)?

Principle 6. Make good use of distinctive cues. Have you created distinctive icons that give the users an unmistakable fix on their location within your site? In contrast, have you made effective use of standard icons that carry their own meaning for the user? Can users find how to enter and leave your pages? Are all functions clearly marked and located where the user expects to find them?

Principle 7. Use the type of cues the users expect for the type of site you are designing. Are your cues reliable, salient, distinctive, and permanent? Do you consistently use them from page to page within your site? Make use of a consistent look and feel for your site. Users may enter your site at almost any page; be sure they recognize the terrain even if they do not exactly land on the page they want to visit. Descriptions of an educational adventure linking USA schools to remote sites is found in 'terra.org' at 'http://terra.org'. The consistent look and feel of this site makes it easy to recognize from any page.

Principle 8. Determine whether your information set is more analogous to a structured or an unstructured physical environment. If the user expects a structured information set, have you used standard formatting within the site? If the user must navigate through an unstructured database, have you facilitated these strategies? An example of a site with an extensive amount of information in a structured environment is Recovered Memories of Sexual Abuse: Scientific Research & Scholarly Resources at 'http://www.jimhopper.com/memory#eb'. This site is basically a tutorial in narrative form of posttraumatic stress disorder. This site allows the user to return to the Table of Contents from each page. It is designed to mimic paper text that is quite familiar to the user. Therefore, the absence of navigational cues is probably appropriate. However, a less structured or a less familiar format would make this site difficult to use and would result in the loss of its valuable information to the WWW user.

Principle 9. Take advantage of the users' natural problem-solving strategies in an unstructured site. Do your pages allow appropriate use of prediction to insure a sense of confidence? If the user does get lost, do you provide features that allow recovery? Do you have a method analogous to catching features that constrain the users' wanderings toward the desired goal? Does navigation within your site allow the user to aim off and still spiral into their goal location?

Principle 10. Try it out under the actual conditions of use. For navigating WWW sites this principle includes two specific conditions: a sample of the users and a sample of the users' computer systems. Can people similar to your target audience find your site and move comfortably around it? Are your design features visible on the screens your users employ? For example, have you designed on a 21 inch screen using a frame that contains critical icons only to find that this frame is not displayed on a 15 inch screen? In this case, do you want to use frames at all?

REFERENCES

Dominguez, C. (1994). Can SA be defined? In M. Vidulich, C. Dominguez, E. Vogel, & G. McMillan (Eds.), *Situation awareness: Papers and annotated bibliography* (pp. 5–15). Dayton, OH: Air Force Materiel Command.

Human engineering guidelines for management information systems. (1985). (Tech. Report No. DOD-HDBK-761). Washington, DC: Department of Defense.

Thorndyke, P. W. (1980). *Performance models for spatial and locational cognition* (Tech. Report No. R-2676-ONR). Washington, DC: Rand.

Tolman, E. C. (1948). Cognitive maps in mice and men. *Psychological Review, 55*, 189–208.

Whitaker, L. A. (1996). Getting around in the natural world. *Ergonomics in Design, 4*, 10–15.

Whitaker, L. A. & Cuqlock-Knopp (1992). Navigation in off-road environments. *Scientific Journal of Orienteering, 8*, 55–71.

Wickens, C. D. (1992). *Engineering psychology and human performance* (2nd Ed.). New York: Harper Collins.

Part II: WEB USER POPULATIONS

6 Children's Online Environments

Allison Druin
Mark Platt
University of New Mexico[1]

THE COMPUTER AS A TOOL FOR LEARNING

Suppose it were music that the nation is concerned about. Our parents are worried their children won't succeed in life unless they are musicians. Our musical test scores are the lowest in the world. After much hue and cry, Congress comes up with a technological solution: By the year 2000 we will put a piano in every classroom! But there are no funds to hire musicians, so we will retrain the existing teachers for two weeks every summer. That should solve the problem! But we know that nothing much will happen here because any musician will tell you, the music is not in the piano! What music there is, is inside of each and everyone of us (Kay, 1996, p. 24).

These are the words of Alan Kay, father of the personal computer. Although his words referred to the integration of music in schools, he was really talking about the integration of technology in schools. He was trying to point out that the mere existence of technology boxes may not be the magical solution for educational problems. Instead, creative ways of using the new tools of technology may help children in their race for knowledge.

Technology can be a powerful tool for problem-solving, developing literacy skills, and supporting interdisciplinary knowledge (Druin & Solomon, 1996; Norton, 1992; Papert, 1985; Spoehr, 1994), but these reasons are not enough to make children want to learn. As most teachers and parents know, children learn when they have a desire to do so. Some say that better grades, a high-paying job, or the novelty of using computers are reasons enough for students to want to learn. We believe, however, that supporting a child's desire for learning can best be done when tools are used to offer an animating idea or an engaging project (Druin, Platt, Grinberg, & Uscher, 1996). Learning for the sake of learning a tool does not offer children rich, complex learning experiences.

---

[1] Allison Druin has relocated to the University of Maryland, College Park and Mark Platt has relocated to the University of Louisiana.

Instead, using these tools to explore places, find information, and communicate with people can lead to powerful learning experiences.

This approach to using technology has been described as problem-centered learning (Norton, 1992). It asks students to explore music, art, mathematics, science, and more by using computers as tools for learning. It puts a real-world problem of interest at the center of the learning process (e.g., Create a web page on local animals; Design a brochure for a school play; Redesign the public library). Students are asked to solve a problem by using the tools of technology as well as the content discipline. Once they have created their projects, students present their work to their classroom peers and to the teacher. Based on the feedback they receive, they may revise their projects. The students are not looking for one right answer, because the problem at the center of their exploration offers many solutions. From this experience, students may not only be motivated to learn, but they may also be able to learn interdisciplinary skills such as art, science, technology or math (Eisner, 1994; Norton, 1992).

An example of a problem-centered use of online technology is the Technical Education Research Center (TERC) Kidsnet curriculum. It offers elementary school students the opportunity to explore such topics as solar energy, nutrition, acid rain, water, weather, and trash by using the tools of technology. The point of their learning is not to think about technology, but to use technology to be scientists. Students are asked to design and build their own rain collectors and to measure the acidity of the local rain water. They are asked to analyze the foods in their lunches for nutrient content. They are asked to build solar collectors and to measure solar radiation levels. These student scientists use online technologies to share information with other classrooms around the world. With other students in Texas, Russia, or Alaska, students have online discussions about the experiments they are working on in their classrooms. They also discuss their work online with a unit scientist, an adult expert in a related field. Bob Tinker, a physics educator, has been the driving force of this educational approach at TERC. He has pointed out:

> The idea that kids are measuring something for others and that the results matter, transforms the classroom. Students become committed, careful, articulate, and surprisingly intense. Along the way they learn large amounts of mathematics, science, technology, experimental procedures, geography, and communication skills in a way that will probably last a lifetime (Tinker, 1993, p. 7).

These students could have conducted some of their activities without the use of online technologies, but their experiences would have been less rich, less complex, and less like what real scientists do. With the use of these tools, students explored, collaborated, conducted research, and published their ideas in ways that were powerful and meaningful learning experiences for them.

THE COMPUTER AS A GATEWAY AND MIRROR FOR LEARNING

Today there is much discussion, excitement, and analysis among researchers, teachers, parents, and even kids about the many possibilities of the Internet. Sherry Turkle, MIT Professor and author of *Life on the Screen*, has written that a computer can be seen as a tool, as a gateway, and as a mirror (Turkle, 1995). What did she mean? According to Turkle, a computer offers society a gateway to more people, places, and information than could ever be thought possible. The gateway she refers to is not a physical place, but a virtual one known as the Internet. It is a technology that has just begun to affect average people's lives. Today, this gateway is taking children, teachers, and parents to people, places, and information that once upon a time was never thought possible.

According to Turkle, this gateway is also a mirror because it can tell us more about who we are as people (Turkle, 1995). And with self-reflection can come questioning and change. Thanks to computer technology, people have begun to question what it means to communicate; what it means to be a community; even what it means to learn. Such educational goals as literacy and problem-solving skills are addressed by the computer as a tool. Educators are just starting to take advantage of the gateway and mirror aspects of computers. If children have the opportunity to use a computer as a gateway to meet new people with new ideas, they may learn about other cultures and practices, and they may learn more about their own cultures and practices. The gateway can offer a mirror for understanding.

For instance, children have traditionally enjoyed discussions with their parents, teachers, and friends in their communities. They have telephoned to distant friends and family and have written to penpals in other countries. With the gateway tools of computers, children have the opportunity for more immediate communication with many more diverse individuals. Not only can they find connections with like-minded people, but they also can be anonymous if they so choose. A 10-year-old boy who is also a 4-year college graduate explained:

> One of the greatest things about the Internet is that no one has to know who you are.... There have been times people my own age may not be interested in what I am interested in, and, at the same time, people who are older than I am, who might be interested, do not always want to spend time with someone so much younger than they are. But if you are on the Internet, things like age are unimportant—or invisible, anyway—if you want them to be (Long, 1994, p. 43).

Another boy said, "I am 15 years old. I came to terms that I was gay last summer and aside from some depression, I am OK. I am not in denial about being gay. I would like to write to someone that I can talk to about issues I cannot talk about with my friends" (Long, 1994, p. 47).

Online environments are important places to connect with other people. They offer very powerful learning opportunities. Our examples have primarily

been based on textual communication, but in the future it will become commonplace to use graphics and sound to exchange ideas in real time online. Imagine that your voice is your own, but the visual character you display to the world is a fanciful dragon, movie character, or comic-book hero. One such experimental system has been developed by the Starbright Foundation in collaboration with software, hardware, and film producers. With this system, children who are hospitalized for serious illnesses can visit children in other hospitals. Yet, instead of wandering through sterile hospital environments, children explore, with others, simulated jungles, caves, and clouds. Steven Spielberg, a creative force behind this project has explained, "For a sick kid, having fun means lots of different things....It means letting the imagination run wild when the body cannot. Or controlling the switches when the pain starts. Or just being taken seriously for a few moments as a kid" (Lewis, 1995, p. C3).

We believe that the power of a computer as a gateway to people, places, and information cannot be overstated. How do computers affect how children learn? We believe that more project-based, interest-oriented, collaborative learning environments are possible, and that new technologies must be created by researchers who understand new approaches to learning. As classrooms change, so too must technologies. Researchers often know little about the needs of child learners. Conversely, children and teachers often know little about emerging technologies that can support powerful new learning opportunities. It is our hope that with books such as this, new discussions, collaborations, and research may be supported in the future.

USING ONLINE ENVIRONMENTS TO SUPPORT CHILDREN'S LEARNING

At the University of New Mexico (UNM), we have begun several research projects that support children's learning with the use of online environments. In the two sections that follow, we describe two examples of research projects that not only inform us about how children use online technologies, but are an ongoing service to the community. We currently support elementary school children in rural and inner-city areas with curricula that focus on music or health education. We support the children with a problem-centered approach to learning and the use of online technology tools. We have seen enormous strides in children's learning of health and music content, as well as in their understanding of the technology they are using. We have used these experiences to inform ourselves as researchers about what new technologies need to be created. The more we work with children, the more we find that they have a lot to tell us about the technology they use.

Example 1: An Online Music Curriculum

In one corner of the room, a group of children experiments with a Chinese musical instrument. In another corner, two children argue about what Irish song to include on their Web page. At the computers, a child is sending E-mail to his

Australian E-mail buddies to ask what music they think he should include on his Web page. This is a typical day in an afterschool program entitled Immersion in Music (Uscher & House, 1995).

This afterschool program was begun in January 1996 at the Hawthorne Elementary School in Albuquerque, New Mexico (Druin, et al., 1996). The school supports a culturally diverse, low-income student population. Many students are children of parents who have recently immigrated from Mexico, Vietnam, China and elsewhere. Most students had little or no access to computers before entering the afterschool program. None had ever used the Internet. In addition, the public schools in Albuquerque had recently cut all support of music and art at the elementary school level. In this environment we chose 10 fourth-grade students (8 to 10 years of age) for an afterschool program supporting the exploration of music by using the Internet.

This program was funded by the Corporation for Public Broadcasting. The funding partially supported four UNM researchers from the Departments of Music, Computer Science, the College of Education, and a representative from KUNM (a local public radio station). These researchers worked with the elementary school's technology specialist to teach the afterschool program. In addition, five Apple Macintosh computers were purchased, and Internet access was installed in the school.

We began the program by working with the children to understand three crucial elements of the program: music, E-mail, and the World Wide Web (WWW). The children listened to recorded music and live performances, and they made their own music with and without the use of the computer. In terms of E-mail, the children discussed what letters were and what questions they might ask people far away from where they lived. To understand the WWW, the children talked about what libraries were and how people published information.

The students were initially confused about the differences between E-mail and the WWW: They saw the two as being the same because both applications were on the same computer screen, both used a modem, and both shared information with other people. To better support their understanding of the differences, we had to have the children perform physical exercises away from the computers. First they wrote a "secret message" on a small piece of paper, then put these messages in balloons and inflated the balloons. We then paired the children and asked them to "send" their secret balloon messages to their buddies. After the balloons stopped flying, the children popped their balloons and read their messages. This process, we explained, was similar to E-mail. The children also drew on a large piece of paper whatever they thought represented what they liked or disliked; we taped these pictures to their bodies. All the children then walked around and looked at each other's pictures. This process, we explained, was similar to the WWW. We talked about the differences between public and private messages, as well as appropriate and inappropriate sharing of information. After these exercises, the children never again seemed confused about the difference between E-mail and the WWW.

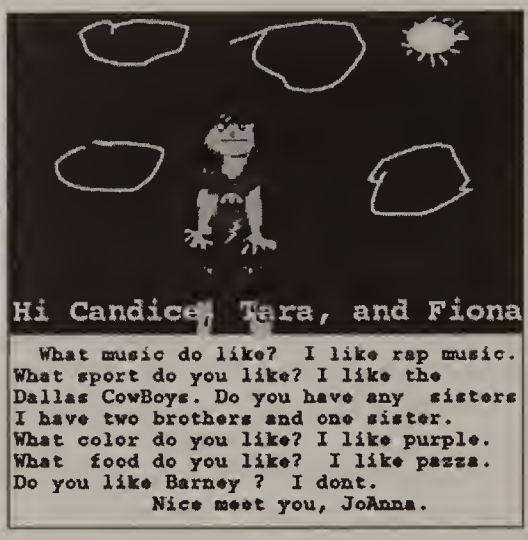


FIG. 6.1. Example of introductory E-mail in the New Mexico afterschool program.

Early in the afterschool program the children established Australian E-mail buddies with a class of children a grade or so older (10 to 12 years old). We adults had made contact with the Australian teacher and organized the teaming of two Australian children for every one New Mexican child. The children in our afterschool program introduced themselves to their E-mail buddies and asked them about the types of music they listened to, their families, even their favorite colors. In response, the Australian children discussed their likes and dislikes and their daily activities. As the program went along, we found that the children used E-mail as a way to reflect on their activities in the program.

As the children's E-mail correspondence continued, the children questioned many things that they had taken for granted. For example, the children in Australia spelled certain words differently from the children in New Mexico. The biggest excitement was over the word *color* versus *colour*. The Australian children were unfamiliar with such things as, rootbeer floats, and one day there was much discussion among the children in New Mexico about the best way to describe a rootbeer float. These are just a few examples of how the children's concepts of language and culture changed thanks to their use of E-mail. Figure 6.1 is an example of one child's E-mail message.

In addition to their online correspondence, the children began to use the WWW. First they used it as a research tool, to look up information about different kinds of music. Then they used it as what they called a music scrapbook. They began collecting links to Web pages that they thought had good information on different kinds of music. Eventually, each chose a partner, and both researched one type of music (e.g., Irish, Chinese, Native American, Australian, and Russian). Each team explored not only the WWW, but also music in books and on records and used musical instruments the researchers shared with them.

Eventually, each team collected enough information to create its own Web pages (House, Platt, Boltman, Druin, & Uscher, 1996). (To see more examples of the children's work: http://kunm.unm.edu/k12/k_12_homepage.html). Figure 6.2 is an example of a Web page created in the afterschool program.

The afterschool program continues today in New Mexico, and more children have been selected for the program. We hope to refine the problem-centered approach to learning and to better understand the conceptual changes that children using online technologies seem to undergo in reference to language, culture, communication, music, and technology. In addition, we hope to learn what new technologies should be created to support educational environments such as this.

Example 2: An Online Health Curriculum

In a rural state such as New Mexico, contact between isolated students, often from minority cultures, and health care professionals is minimal; indeed in many communities students and health care professionals meet only when the students need medical attention. We have established an electronic connection between the UNM's Health Sciences Center and a New Mexico elementary school system which enables students of all ages to discuss matters of medical interest and more

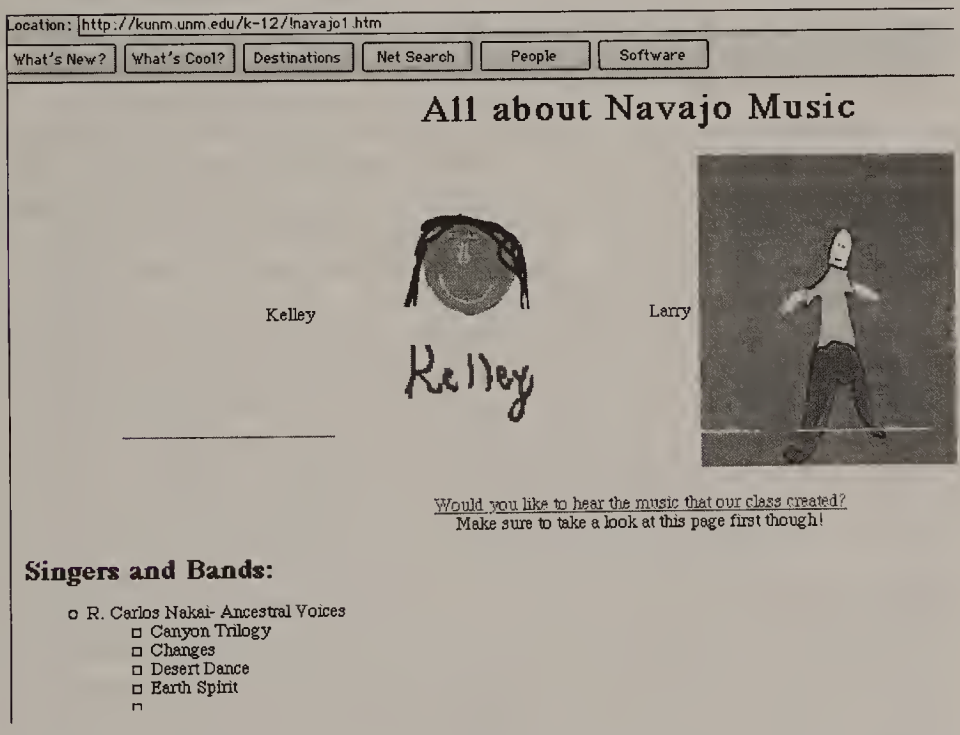


FIG. 6.2. A part of an example Web Page created in the New Mexico afterschool program.

specifically to be taught a health curriculum via computer. Our purpose is to provide answers to students' health questions and support for teachers dealing with increasingly complex health problems in a world where preventable diseases are the number one cause of morbidity and mortality.

Program Aims and Design. The program for connecting students to the UNM's School of Medicine had four main aims: to provide current health information to school students; to answer students questions about health and human biology; to enable minority and women students to contact health professionals who provide career information; and to teach children how to communicate with distant experts.

For a communication project to succeed, an elementary school must have enough networked computers to enable easy student access to check E-mail and to post new questions. Although the WWW offers connectivity and information-gathering abilities to school students, we designed this program to work over a simple E-mail system. The system we chose was FirstClass by SoftArc, which in the simplest configuration, needs only a school network and outside access either from a modem or via the Internet. The FirstClass E-mail application is icon-based, and is designed primarily for grade school students. Although it does not have a spell checker, it does allow the use of different fonts and colored text in the message, and enables users to attach and download any file. For students in kindergarten and first-grade classes to use the system, we encouraged the teachers to enter the questions, and then communicate the replies to the children.

At UNM's School of Medicine, all incoming students must buy personal computers, and all members of the faculty have access to a networked computer. The medical students have both a network and dial-in access to a local FirstClass mail system so that distribution of questions from the elementary school students to the medical school faculty and students was achieved quickly and efficiently.

Health Curriculum Design. Students in Grades 2 to 5 were able to achieve direct access to health care professionals by sending specific or general health-related questions to a personal computer user known as Dr. Clean and Healthy. This address was checked at the end of each school day for questions by a medical professional who acted as a gatekeeper. The gatekeeper would make decisions about forwarding the questions to physicians or volunteer medical students. Medical students were included in the program to train them in the phrasing of age-appropriate answers. After the questions were forwarded, the participants were instructed to return the answers in 12 hours. We had determined that a rapid answer to the school students' questions was of critical importance in maintaining the children's interest in the project. The gatekeeper forwarded the replies to the correct addresses after first checking the applicability of the answers. Technical terms (e.g., DNA, meiosis) were presented in bold text to allow teachers to provide background information if needed. Answers were complemented by further questions to stimulate class discussion. Students'

answers developed into "threads" similar to those encountered on Usenet discussion groups:

Q: I just wanted to know why do some people have different color skin. Our class has been studying Black American history. Adam
A: Adam, we all have a color or pigment in our skin. It's called melanin. Black people have much more of this than Caucasian or White people. Some scientists think that having dark skin was an advantage in a hot and sunny continent like Africa so that there was a selection for people with dark skin. We do know that the melanin protects the skin from burning because it removes much of the dangerous ultraviolet (UV) rays. Black-skinned people have fewer skin cancers than light-skinned people. If you would like to know more about what a "selection for people with dark skin" means, then ask—Dr. Clean and Healthy.

Q: How come people get cancer from the sun? Donna, 3rd grade
A: Dear Donna, This is a good question, especially in sunny New Mexico. Sunshine contains a certain kind of energy called ultraviolet radiation. This radiation hits your skin and causes damage. Most of the time this damage is seen as a "sun-burn" where your skin stings and gets real red and itchy. Many scientists think that if you get too many of these sun-burns, your skin eventually has a hard time repairing itself. In a few people the "skin cells" which are not repaired are no longer normal. They can grow out of control, because the body doesn't know how to get rid of these cells. These growing out of control cells are skin cancers.

NOW I HAVE A QUESTION FOR YOU!
Why should you wear sun-screen?

Other questions from first-grade students included:

1. When we have just got over a cold, why won't the coughs go away?
2. Why can't we sleep even when we're very tired?
3. Why do we have nightmares, and why do we feel like bugs are crawling on us when we're trying to sleep?
4. Why do we get bumps on our skin when we run?
5. Why do spider bites itch?
6. I have a white bump on my tongue (not a taste bud) that hurts. What could it be?
7. Why do our noses get sniffly when we don't really have colds?
8. Why do we talk in our sleep?
9. Why do we get cramps when we run?
10. Why is it hard to breathe deeply (catch our breath) after we've been running?
11. Why is it hard to talk after we've had a coughing attack?

Some questions were of educational value not only to elementary school students but also to medical school students. The question about talking when asleep occupied the medical students for some time; now the elementary students understand that there are scientific problems that do not present clear solutions.

In another project, the relationship between diet, exercise and other health habits was considered. Students electronically received a spreadsheet on which to enter their daily food intake, exercise, and so on. After data entry, the students measured their own height (m) and weight (kg). A class spreadsheet was then returned electronically to the medical school, and the results were compiled in a graphic format with general and, if required, personal recommendations. This exercise had many direct benefits: The elementary school students gained basic knowledge about the relations between diet, exercise, and weight, and they practiced the metric system and learned how to manipulate data in a spreadsheet.

Children's online environments that support distance learning can achieve many educational goals. The elementary school students were able to ask questions about any subject of interest while gaining factual knowledge and computer communication and data manipulation skills. The medical school participants gained valuable experience in framing age-relevant answers to common questions and, in many cases, had to research answers themselves.

WHAT ABOUT THE FUTURE?

Our work with children has led us to conclude that three areas of focus should be considered in developing online environments for children in the future. The first area is the need for more flexible online tools for children. Most children do not like to type as much as they like to draw. If there were easier tools to enable children to send their drawings via E-mail, our children would have been thrilled. We had to use three different applications to send drawings via E-mail and had to hope that the children on the other end had the applications needed to see what we had done. In addition, a spell checker for the children's E-mail program would be a welcome addition; though they could have used a word processor and then copied and pasted the results into their E-mail program, the tools today are more difficult than they need to be.

Another area for future development is that of online collaboration tools. The collaboration we refer to is not with people far away but with those in the same room. Children like to work together on one computer, but few if any, tools support the work of two children using one computer.

The last area for future development that we see as important is the area of improved visualizations of online environments. Our music students' confusion between E-mail and the WWW may have been due to the fact that there is little difference between the way E-mail and the WWW look on a computer screen. On the WWW, few programs show children in a visual way what pages connect to other pages and how many people are looking at the same pages they are. On the

form of communication different from the WWW. New types of visual cues would be invaluable in supporting children's understanding of different forms of online communication.

Today, there seem to be few fully integrated online environments, developmentally appropriate for elementary school children. Children have told us what they want and what they don't want. We believe that there are many possibilities for exciting research in this area of online environments for children. We look forward to what the future may bring.

REFERENCES

Druin, A., Platt, M., Grinberg, J. G. A., & Uscher, N. (1996). Pedagogy, content knowledge, technological literacy, culture and equity: The cart leading the horse. In the *Proceedings of the American Association for Teaching and Curriculum* (pp. 12–13), San Antonio, TX.

Druin, A., & Solomon, C. (1996). *Designing multimedia environments for children: Computers, creativity, and kids*. New York: Wiley.

Eisner, E. W. (1994). *Cognition and curriculum reconsidered* (2nd ed.). New York: Teachers College Press.

House, D., Platt, M., Boltman, A., Druin, A., & Uscher, N. (1996, May). Immersion in music: Connecting children and cultures through music. http://kunm.unm.edu/k-12/k_12_homepage.html (Last accessed October 1996).

Kay, A. (1996). Revealing the elephant: The use and misuse of computers in education. *Educom Review, 31*(4), 22–28.

Lewis, P. (1995, June 8). Ill children to get 3-D playground out in cyberspace. *The New York Times*, p. C3.

Long, M. (1994). We are the world. In *Netguide* (pp. 43–47). New York: CMP Media.

Norton, P. (1992). When technology meets subject-matter disciplines in education: Part three: Incorporating the computer as method. *Educational Technology, 32*(8), 35–44.

Papert, S. (1985). Computer criticism vs. technocratic thinking. In the *Proceedings of Logo'85: Theoretical papers* (pp. 53–51). Cambridge, MA: MIT Press.

Spoehr, K. T. (1994). Enhancing the acquisition of conceptual structures through hypermedia. In K. McGilly (Ed.), *Classroom lessons: Integrating cognitive theory and classroom practice* (pp.75–101). Cambridge, MA: MIT Press.

Tinker, R. F. (1993). Telecommuting as a progressive force in education. In *Technical Education Research Center Technical Report*. Cambridge, MA: TERC Publications.

Turkle, S. (1995). *Life on the screen*. New York: Simon & Schuster.

Uscher, N., & House, D. (1995). *Immersion in music: UNM/public elementary schools partnership, Global perspectives through the Internet*. Grant funded by the Corporation for Public Broadcasting, University of New Mexico, and KUNM Radio.

7 Designing Web Pages and Applications for People With Disabilities

Lila Laux
U S WEST Communications

Users and designers of Web pages have typically been young, educated, computer literate, and able-bodied. However, when designing a page or application for the World Wide Web (WWW), designers should keep in mind that the target population includes millions of potential users of Web pages who have various handicapping sensory and physical conditions. It is also important to remember that the U.S. and world populations are rapidly growing older, so the number of people with such handicapping will significantly increase in the next decade. People with handicapping conditions represent a huge user market for Web applications; not only do they want to use the Web, but their families, teachers, and employers also want them to be able to use it. If you, as a designer, choose to design a Web site that is not accessible to people with handicaps, that is certainly your right. But Web page designers should be sure that it is their choice to create an inaccessible Web site, not simply an oversight or an error, and they should consider the consequences.

WHY WORRY ABOUT ACCESSIBLE WEB SITES?

Over 30 million people in the United States (Vanderheiden, 1996) and many more millions worldwide have physical, sensory, or cognitive limitations that make interacting with traditional monitor, keyboard, and mouse configurations difficult. Assistive and adaptive computer technology can enable most people with disabilities to use a computer efficiently and effectively (Lazzarro, 1993; 1996). Adaptive technology today supports almost all users with special input or output needs. People with disabling conditions also want and need access to the WWW; but, because of the design of many Web pages, full access is not possible even with high-end adaptive technology.

The Internet and the WWW offer people with disabilities opportunities that were only dreamed of a few years ago. Information access, online communication, and interactions with complex systems are now possible for users who are unable to move, see, speak, or hear.

In this chapter, little space is devoted to discussing the adaptive devices for computer systems that allow people with a handicapping condition to get access to the Internet and the WWW. This technology is evolving rapidly, and any

descriptions of what is available would soon be out of date. Such technology is available, and qualified specialists can help users select the assistive devices they need to be able to interact with a computer and with the WWW (Lazzarro, 1996).

But, although being able to interact with a computer is a necessary prerequisite to using the WWW, that is not the same as interacting with a Web site. Even users who have access to a computer and can access the WWW, do not necessarily have effective and efficient, or even acceptable, access to information or functionality on Web sites.

The Right to Accessibility

Web page and application designers might believe that they have the right to design a Web page to meet their own needs or aesthetic goals, but today many Web sites are used by workers or teachers. When a Web site is used in a job or in schools or universities, accessibility becomes an issue that may be addressed by the Americans with Disabilities Act of 1990 (ADA).

When employers or educational institutions commission or develop Web Sites for an intranet or the Internet, they should make every effort to see that they are accessible. If employees need to use an outside Web site for a critical job function, the employer or institution may be responsible for providing adequate access.

If Web sites are designed so that current adaptive equipment cannot make the pages accessible, employers and educational institutions may have difficulty providing acceptable accommodations. For example, if a page is heavily loaded with graphics that are not alt tagged or tables that cannot be read effectively by a screen reader, an employer or educational institution could conceivably be required to hire a sighted person to read and interpret the page for a visually impaired employee or student who needs access to perform required work. Similarly, a service provided to the public via a Web site or page that is not accessible to users with disabilities may be subject to an ADA claim; and, more important, the service provider may lose market share because many potential customers are unable to access the services. These examples are not meant to scare Web designers with threats of ADA claims. Easy solutions can prevent these scenarios.

The WWW and Handicap Access

This chapter is intended to alert Web page designers and developers to the issues related to access by people with handicapping conditions. It is not a detailed "how-to" for so-called electronic curb cuts (Murphy, 1996). Technology related to Web applications is changing rapidly and information about what is available becomes rapidly out of date. Numerous Web sites provide up-to-date information and comprehensive guidelines for developing and testing Web sites for accessibility.

One of the best sites providing "how-to" information for designing accessible Web sites is supported by the Trace Research and Development Center at the University of Wisconsin (http://trace.wisc.edu/). The ACT Center, which is

sponsored by the Public Service Commission of Canada (http://www.igs.net/~starling/acgen.htm), and the World Wide Web Consortium (http://www.w3.org/pub/WWW/Disabilities/) also offer extensive and up-to-date information and guidelines for designing accessible Web pages.

Instead of presenting a detailed template for creating accessible Web designs, this chapter presents the concept of universal design (Connell et al., 1995) for Web pages and Web applications. Universal design refers to the design of products and environments that are usable by all people, to the greatest extent possible, without the need for case-by-case accommodation. If you adopt universal design when developing a Web page, your pages will be more readily accessible to most people with disabilities who are already using computers.

Basic Accessibility Issues for Web Sites

After a user with a handicapping condition has access to the Web (adapted computer system, modem, account, and so on), there are still three basic types of problems associated with making the information in Web sites accessible and allowing users to interact with Web sites. These problems include making printed text and graphic information available to users with vision problems, making audio information available to users with hearing impairments, and providing an interaction method for users with visual or motor problems that affect their ability to use a mouse or keyboard.

The following discussion covers universal design issues as they relate to various handicapping conditions. Not all handicapping conditions are addressed; instead, this brief discussion focuses on the access and functional problems that confront the largest groups of users with handicapping conditions who want to use the WWW. These groups are people who are blind or visually impaired, deaf or hard-of-hearing, or have motor problems relating to input devices. An extensive discussion of disability incidence and barriers to WWW use can be found on the Trace Research and Development Web site at http://trace.wisc.edu/text/univdesn/populat/populat.html.

Many people with a handicapping condition have multiple handicaps. For example, people who have lost arms or hands in accidents often have sustained some brain damage or had a closed head injury resulting in spasticity. People with multiple sclerosis may simultaneously experience vision, motor, and memory problems, and their problems may fluctuate significantly over a brief time. In addition, most people experience some combination of vision, hearing, and musculo-skeletal or motor problems as they become older.

General Guidelines for Making Web Pages Accessible from DO-IT

Disabilities, Opportunities, Internetworking, and Technology (DO-IT) at the University of Washington (http://weber.u.washington.edu/~doit/) is a federally funded program to increase the participation of individuals with disabilities in science, engineering, and mathematics academic programs and careers. One of the services that DO-IT provides is to present workshops about adaptive technology.

DO-IT offers the following general guidelines for making Web pages accessible to the broadest range of users with disabling conditions:

1. Use a consistent, simple layout across all pages of the site.
2. Ensure that contrast between the background and foreground (especially text or other symbols) is high.
3. Provide a text alternative to all graphics and image maps (alt-text or D-tag).
4. Provide descriptive captions for all pictures.
5. Use tables and frames sparingly or offer alternative pages for screen readers.
6. Use large buttons and controls.
7. Caption video and transcribe audio displays.
8. Insure that your information is compatible with screen-reading software and other adaptive technology.

MAKING INFORMATION AVAILABLE TO USERS WITH VISUAL HANDICAPS

According to the American Federation of the Blind, among the approximately 9 million people in the United States who are blind or have low-vision (Elkind, 1990), there are about 900,000 blind or low-vision computer users (approximately 10%). The percentage of those 900,000 people who also use the Internet is not known, but the National Library Service for the Blind and Handicapped at the Library of Congress has estimated that only about 1% of the 20,000 to 30,000 blind people in the Washington, DC, area currently use the Internet (Tannenbaum, 1996).

For blind and low-vision computer users who also access the Internet, graphics-intensive Web sites pose a serious challenge: most blind and low vision users use text-to-speech screen readers. Text-to-speech screen readers often do not work well in Windows environments or with any applications which employ graphical user interfaces (GUIs).

Most GUIs utilize mouse-based or point-and-click navigation, rather than keyboard-based navigation. Blind and visually impaired users have difficulty using a mouse to navigate on a GUI (they cannot locate the pointer on the screen). But simply providing a keyboard-based browser does not address all the navigation problems that occur when hypertext links on Web pages are entirely graphic, or when there are many links on a single site.

What is the recommended design for Web users with visual handicaps? Always design to be screen reader friendly. Screen readers convert text on a computer screen into speech. Blind and visually impaired computer users employ screen readers to access information on a screen and typically use keyboard commands to navigate through the screen.

Designers who want their sites to be accessible to those who use screen readers must remember the following problems often associated with Web pages and screen readers:

1. Screen readers read an entire line of text. If there are many links on the same line, it is difficult to tell which link a cursor is pointing to. Links embedded in paragraphs are hard to navigate because it is not apparent to users where the link is. If there are frames, the screen reader does not always recognize them and skips across the boundaries in reading. But frames can be made more accessible to screen readers by following these basic guidelines:
   A. Use the noframes option for users employing text readers like Lynx.

   B. If a homepage has frames, consider creating an introductory page before the home page and offering access to the frames or noframes version of the site.
   C. Create individual pages for each frame on a page and place links to the individual frame pages on the page that links between frames. Those who use a graphical browser have the option to use the frames or go to the individual pages.
   D. Be sure that scrollbar controls on each frame are labeled so that screen readers render them intelligibly. Users may lose information in a scrollable frame if they are unable to scroll because they do not know about or can not find the controls in the slider.
2. Graphic images and icons that are not text (alt) tagged are indecipherable to screen readers. Links and graphic images do not interfere with screen readers, but if designers do not supply an alt or text tag for them, the screen reader only sees (and says) [LINK], [INLINE], or [IMAGE]. Pages made by Web page-creating applications, which produce the Hypertext Markup Language (HTML) to present the page you design, do not include alt tags. Unless a Web page developer knows how to go into the HTML created by these packages and add alt tags, the images and icons on the pages created by these programs will not be accessible to screen readers.
3. ISMAPs and BISMAPs are features that are entirely graphic and entirely inaccessible to the screen reader. Designers with a totally graphic page, must provide comprehensive alt tags or an alternative text-only page for screen readers.
4. Forms are difficult for screen readers to translate and practically unusable by users who are visually impaired.

Developers should always view their sites with a text-only browser. This is the ultimate usability test. If a site cannot be understood when displayed by a text-only browser, it is essentially inaccessible to vision-impaired Web users. (To

view a site without the graphics and to see the alt tags that the screen reader will read, go to the Options menu and turn off Auto Load Images before loading the page.) Or, developers can view pages with the widely used text browser, Lynx. To do this, go to the form at http://www.miranova.com/~steve/Lynx-View.html and enter the Uniform Resource Locator (URL) of the page. Developers can also have a site reviewed (free of charge) for HTML accuracy and accessibility at http://www.htmlhelp.com/links/validators.htm.

Examples of bad HTML (HTML that is unintelligible to screen readers) and good HTML (corrected so it can be read intelligibly by screen readers) appear in *Designing Open Communications: An Insurance Policy for All* (Wakefield, 1996). Other examples of incorrect and correct HTML are found on Gregory Rosmaita's Web pages at http://www.hicom.net/~oedipus and http://www.hicom.net/~oedipus/blist.html. Rosmaita, a blind Webmaster, designed these Web sites himself and made the following suggestions for "quick and simple" changes that make pages friendly to a much wider range of browsers:

1. Make sure icons and images do not show up as simply [INLINE] by adding an alt tag to <img src="icon.gif">. Alt tags tell the function of the graphic.
2. Depending on the type of image being displayed, use one of the following alt tags: <img src="icon.gif" alt="[Give description - include type GIF/JPG and size in KB]"> for a large icon or picture.
3. More examples for using alt tags: for a bullet: alt="*"; for a line: alt="----------"; for a logo: alt="Name that logo stands for"; if nothing should be spoken: alt=""; for a capital letter in a word: alt="W">ashington; when using a counter that outputs the number as an image: alt="next number."
4. To prevent the text-only browser (and therefore, also the screen reader) from rendering simply [LINK], add an alt tag to the image in the link: for example, instead of <img src="mailus.gif"> (which displays [LINK] on a text-only browser), add alt="Send us E-mail" after the img src. Now users see (and hear if a screen reader is employed) Send us E-mail instead of [LINK] - much more usable (Rosmaita, 1996).

If there is more information than can easily be put into an alt text, also include a D-tag or description anchor, which is an upper case d or a short description, that acts as an anchor to a text description on another page. Gregg Vanderheiden, Director of the Trace Center, has stated that such description anchors should gradually become less necessary, because the new graphic standard called PNG allows a text description to be stored directly in a picture's data format. But he recommends that until everyone adopts this format, the D-tag should be employed whenever a blind person or person without graphics capabilities would miss information not otherwise presented on a page (Vanderheiden, 1996).

Screen Magnifiers

Some visually impaired users employ screen magnification programs to enlarge information on the screen to a size that they can see. Screen magnification programs, such as *IBM ScreenMagnifier/2* (IBM Corporation), *LPDOS* (Optelec US Inc), *ZoomText* (Ai Squared), and *Magic* and *Magic Deluxe* (Microsystems Software, Inc.) magnify screen elements up to 32 times. (For more information about these and other adaptive products and their manufacturers, go to ABLEDATA at http://www.macroint.com/cgi-win/ableweb.exe?prc=int&q=sq.) These magnification programs do well with text but may render large tables and most large graphics or pictures incomprehensible.

Using Audio as an Alternative to Descriptive Text

When using audio clips on a Web page, be sure to indicate their format (*e.g.* .WAV or .AU) and file size in kilobytes to prevent people from downloading a huge file and then discovering that they can not play it because it is too big. The new audio-on-demand applications make it easy to provide audio information to users with low vision. These applications deliver a highly compressed stream of audio information to a computer in real time. Very long clips can be delivered via these systems more efficiently than by downloading .WAV or other digitized audio formats. Many users do not want to download a million byte file just to hear a few moments of speech or music, and audio-on-demand applications may solve this problem, although the quality of sound may not be the best. A major problem with audio-on-demand is the need for special software on a Web server to deliver it, as well as local audio input and output capabilities. To deliver audio-on-demand from a site, tell site visitors how to get the client software necessary to receive the format.

As multimedia personal computers become more common and as audio-on-demand systems become better and more widely accepted, real-time audio may become a viable alternative format for describing graphics or even delivering the entire content of a Web site to users with visual handicaps. But using audio-on-demand without also providing text information will still exclude people with auditory handicaps.

Braille Displays for Deaf and Blind Users

Braille translators for computers are also available to assist users with visual handicaps. These devices read screens and translate their contents into a refreshable Braille image created by pins which are raised to create Braille text. Although only about 10% of blind users are proficient at Braille, Braille translators are valuable for Web users who have both visual and auditory handicaps. Although this population is small, keeping both vision- and hearing-impaired users in mind when designing Web pages, makes them more accessible to this small group of users.

Researchers at Georgia Institute of Technology, funded by The Toyota USA Foundation, have discovered that tactile image size is an important variable in making sites available to people using a refreshable Braille display. Their

studies showed that "images larger than the size of an open hand are difficult to explore and comprehend." They chose an image size of 2½ by 2½ inches because "at this size, exploration is quick, even with just one finger", and some preliminary tests indicated that people have little problem "putting the image together" in their minds as they explore. For more information see "Tactile Image Techniques" on the Georgia Institute of Technology Web site at http://dots.physics.orst.edu/tactile/tactile.html.

Hearing Impaired Users

How can developers make the audio clips they include on their sites accessible to users who are deaf or hard of hearing? To make the sound information accessible, developers must provide text alternatives for hard-of-hearing users. If a page contains a speech or spoken passage, provide a link to a text version. (People who are deaf or hard of hearing need this accommodation, and it will also be appreciated by the many Web users who do not have sound-cards in their computers.) Developers need not provide text versions of music or sounds meant mainly for aesthetic site enhancement, but they may want to provide text versions of the lyrics of songs (with the copyright holder's permission) or a description of sounds (squeaking door sound, cat meowing sound), if the sounds make the site more interesting or enhance the utility of the site in some other way (Public Service Commission of Canada, 1996).

Users With Physical and Motor Limitations

For people with poor motor control, such as is associated with cerebral palsy, Parkinson's disease, or upper limb paralysis, keyboard navigation is generally preferable to point and click navigation. Users with motor problems may use adaptations such as key guards or "slow" keyboards that do not accept input as frequently as regular keyboards. Many of these users are unable to use a mouse or a mouse-like device such as a track ball for input. For these users, it is critical that a site be navigable from the keyboard and that keyboard navigation provide access to the various links and other controls. If keyboard navigation is impossible, or too cumbersome, users are unable to take advantage of all the information on a page or to use a page as intended.

CONCLUSION

Web pages are created to provide information or services to people who use computers. Many computer users and potential computer users have a handicapping condition with regard to computer use. There are legal, economic, and ethical reasons to make Web pages available to everyone who wants to use them, including people with disabilities. It is technically possible, and not unduly difficult, to assure that most people who want to access and use your Web pages can do so. It seems reasonable to expect every Web developer and designer take

into account the needs of people with handicapping conditions in designing and developing their Web pages.

Countless resources on the WWW guide designers and developers in creating Web pages that are accessible to most users. As Rosmaita noted, the original idea of the WWW was to design the cores of the pages for the beginning users and "spiffy" them up for more advanced users (Rosmaita, 1996). Designers must keep this guideline in mind as they design Web pages. Sites can be evaluated to assure their accessibility using The Web Page Accessibility Self-Evaluation Test (Public Service Commission of Canada, Diversity Management Directorate, 1996).

REFERENCES

Connell, B. R., Jones, M., Mace, R., Mueller, J., Mullick, A., Ostroff, E., Sanford, E., Story, M, & Vanderheiden, G. (1995, December 7). *The principles of universal design.* http://trace.wisc.edu/text/univdesn/ud_princ/ud_princ.html.

Elkind, J. (1990). The incidence of disabilities in the United States. *Journal of the Human Factors Society, 32*(4), 397–406.

Lazzarro, J. (1993). *Adaptive technologies for learning and work environments*. Chicago: American Library Association.

Lazzarro, J. (1996). *Adapting PCs for disabilities*. Reading, MA: Addison-Wesley.

Murphy, A. R. (1996). *Design considerations: Readers with visual impairments*. http:// www.lcc.gatech.edu/gallery/dzine/access/.

Public Service Commission of Canada (1997, March 2). *Making audio accessible.* The Accessible Computer Technology Centre. http://www.starlingweb.com/ acc/acaud.htm.

Public Service Commission of Canada Diversity Management Directorate (1997, March 27*). Web page accessibility self-evaluation test.* http://www.psc-cfp.gc.ca/dmd/access /testver1.htm.

Rosmaita, G. (1996, November 21). A plea for Lynx friendliness. Rosmaita@ email.njin.net.

Tannenbaum, D. (1996, October 3). *Blind put faith in the Internet.* http://www.Washingtonpost.com.

Vanderheiden, G. C. (1996, October 10). *Thirty something (million): Should they be exceptions?* Trace Research and Development Center, Waisman Center and Department of Industrial Engineering, University of Wisconsin-Madison. http://trace.wisc.edu/.

Vanderheiden, G. C. (1996, November 11). RE>Web standards question. Uaccess-l list email. po@trace.wisconsin.edu.

Wakefield, D. (1994, September 9). *Designing open communications: An insurance policy for all.* The Clearinghouse On Computer Accommodation, General Services Administration. http://www.gsa.gov:80/coca/online.htm.

8 The World Wide Web as a Teaching Resource

Peter Burden
Jenny Davies
University of Wolverhampton

HISTORY AND CONTEXT

The first computer was connected to the Internet in 1969, and in subsequent decades, computer networks grew slowly around the world. In the United Kingdom (UK), computer networks for academic use are now provided by the United Kingdom Educational and Research Networking Association (UKERNA), which operates, among other things, the Joint Academic Network (JANET). JANET is funded by "top-slicing" the government funding for education so that individual institutions do not have to pay for network connectivity, except to the extent that their government funding for other purposes is reduced. When JANET was first set up in the 1980s, it used its own set of protocols related to the ISO/CCITT protocols. By the early 1990s, however, it became obvious that the use of Internet protocols would soon be universal, and JANET was switched to Internet protocols. By 1995, an increasing number of local colleges and schools were connected to JANET and using Internet protocols. JANET is connected to other U. K. Internet access providers via the London Internet Exchange, but UKERNA provides links to North America and Europe exclusively for the use of institutions connected to JANET.

Wolverhampton is an industrial town of about 250,000 people in the center of England, on the western edge of the West Midlands industrial conurbation. The University of Wolverhampton was founded as Wolverhampton Technical College in 1932 and has gradually expanded since becoming a polytechnic (a degree-issuing institution without the title of University) in 1969 and a university in 1992. Its first computer science degree course ran in 1967, when the main equipment consisted of an ancient relay computer built by the U. K. Atomic Energy Research establishment in 1948 and donated in 1957. This machine, which can still be seen in Birmingham Science Museum, had 90 memory locations and took 2 seconds to add up two numbers. The present equipment consists of well over 2,000 personal computers mostly in general purpose, open-access laboratories, along with various Sun (Sun Microsystems Inc., http://www.sun.com) servers and specialist workstations.

In 1984, the then Polytechnic acquired its first off-site computer network connection using the U. K. PSS service for E-mail only, on a dial-up basis. In 1988, a permanent 9600 baud JANET connection was obtained, which used JANET's own protocols. In October 1992, the connection was upgraded to a 64 kbaud line using ISDN and Internet protocols. The connection was quickly saturated, and in July 1995 it was upgraded to 4 Mbaud. The university provides several (currently about 20) dial-up lines for student use running at a maximum of 33.6 kbaud. All computers are networked and Internet enabled.

Readers outside the U. K. should know that at the time of writing (November 1996), a typical high performance personal computer cost approximately £1,100, about 2 to 3% of the average professional salary. Telephone calls were charged by the minute; a local call typically cost about £0.02 per minute (off peak). A number of the university's modems are connected to the telephone system operated by the local cable television company, and calls to these modems are free (off peak) to customers who also buy a cable television package. This package is very popular with students. The nearly 200 Internet access providers in the U. K. usually charge around £10 per month on top of which users pay local telephone call charges.

Use of the Internet for Teaching

In the University of Wolverhampton's School of Computing, several lecturers have for some time been interested in using the Internet to support teaching activities. We have particularly concentrated on a course called Social, Legal and Professional Aspects of Computing known as SLAPA throughout this paper. This course contributes 12.5% of the credits for the final year of the computer science degree programs; it is also included in various postgraduate and foreign exchange programs. The course has been a good choice to study the use of the Internet: The students have a wide range of computer backgrounds and skills. The SLAPA course has used the ubiquitous public domain Internet tools (Dern, 1994), anonymous file transfer protocol (FTP), Gopher, and most recently the World Wide Web (WWW) to deliver resources and pointers to students.

Anonymous FTP

FTP is the Internet mechanism for transferring files from one computer to another. People normally require a log-in account on the remote computer before transferring a file from it, but the many Unix systems have a special user account known as anonymous which does not require a password and thus allows anybody to access files. This process effectively allows people to publish material via the Internet by placing it in the anonymous file area.

As soon as the university was connected to the Internet, we discovered that many files were available at no cost from any one of a significant number of anonymous FTP servers. We also discovered that several Archie servers maintained searchable indexes to the material available on the various anonymous FTP servers across the Internet. For our computer scientists, this was a bonanza,

but much of the material was computer science related, and anonymous FTP sites seemed to offer little to support general educational needs.

This lack of support was partly due to the difficulty of identifying resources by file name (Archie servers identify only file names), particularly when the resources are subject to the MS-DOS 8 + 3 constraint, and partly due to the unfamiliarity with the Internet among university staff. The computer scientists persisted and helped by the discovery of a spare disk drive, constructed the university's own anonymous FTP archive, which serves material to a few specialist courses such as Business Data Communications and makes staff-authored material publicly available. The creation of an anonymous FTP archive is not difficult, particularly for those familiar with Unix, but the standard software available had deficiencies that we had to correct ourselves.

Here is an extract from the results of an Archie search for privacy:

```
Host rzbsdi01.uni-trier.de

 Location: /pub/info/faq
 DIRECTORY drwxr-xr-x 512 Mar 20 1996 alt.privacy
 DIRECTORY drwxr-xr-x 512 Feb 18 1996 comp.society.privacy

Host ftp.seas.gwu.edu

 Location: /pub/rtfm/alt
 DIRECTORY drwxr-xr-x 512 Sep 27 04:08 privacy

Host ftp.rice.edu

 Location: /tmp_mnt/ftp/ftp/net/cwis/cwis/gopher/world/Policies/.old/Owlnet/.cap
  FILE -rw-r--r-- 46 Jan 1 1970 privacy
 Location: /tmp_mnt/ftp/ftp/net/cwis/cwis/gopher/world/Policies/.old/Owlnet
  FILE -rw-r--r-- 162 Jan 1 1970 privacy

Host ftp.cs.uchicago.edu

 Location: /pub/discussions/hypernews/cpsr
  FILE -rw-rw-r-- 456 Aug 11 1995 privacy
```

As can be seen, apart from the file names, this extract gives little clue to the nature and format of the information; teachers intending to locate resources in this way must spend a considerable amount of time and effort downloading and reviewing resources to determine their suitability.

Gopher

Gopher is a menu-driven system of linked Internet resources. It enjoyed considerable popularity in the early 1990s before the general acceptance of the WWW. In 1994, we decided to use Gopher to deliver resources related to computer privacy to students studying the SLAPA course. Once a student had started the Gopher application, he or she was greeted with a menu looking like this

```
Internet Gopher Information Client v2.0.16

 Computer Privacy Documents for SLAPA

 --> 1. Jackboots-on-the-Infobahn----John-Perry-Barlow
 2. Jacking-in-from-a-Non-Government-Approv
 3. Jacking-in-from-the-Congressional-Port-
 4. Jacking-in-from-the-Financial-Port----Brock-Meeks
 5. Letter-from-WIRED-Online
 6. NIIagenda.txt
 7. The-End-of-Privacy----Brock-Meeks

Press ? for Help, q to Quit, u to go up a menu Page: 1/1
```

The cursor key can be used to move the pointer up and down, and the Return key to select a particular item. All the items shown here had been found on the Internet by various means and were copied to our local Gopher server. The descriptive text shown on the previous menu corresponds to the file names on the server. It was also necessary to create a Capability directory entry.

Creating the necessary files in the Gopher hierarchy was not totally straightforward; it required some system knowledge. If we had persevered with this approach, it would doubtless have become simpler. The students seemed happy with this way of providing information, and regarded it as a significant improvement on using Archie and anonymous FTP archives. The main disadvantage was the fact that these were local copies of text files and did not contain pointers to any further resources; thus teachers were responsible for selecting suitable files. The practice of making local copies is also questionable in the context of intellectual property rights, but our slow link to the Internet at the time forced us to take this approach.

Although Gopher has now been eclipsed by the WWW we still find it a useful means of navigating the large number of text and other files in our anonymous FTP archive, particularly for users accessing the archive via the WWW. We intend to keep the Gopher server operational for this purpose.

The World Wide Web

The development of the WWW started in 1990, but it was not used at the university until late 1993, and the first server was not set up until mid-1994. It quickly became apparent that the WWW was an effective way to distribute attractive, prepared material to a wide audience. The ability to include links to other resources made creating a list of resources for students simple, once a teacher had identified the resources.

The SLAPA module was one of the first to adopt this approach. The WWW was used to make lists of resources available to students, and students were also required to prepare Web pages on selected topics as part of the coursework assessment of the module. As we have reported elsewhere (Burden & Davies, 1996), this successful although comparatively slow access to some resources resulted in some unfavorable comparisons to resources delivered via conventional library methods and via CD-ROMs. It was pleasing to note that students were happy to prepare and present the results of their own work as Web pages.

The university's Internet link was upgraded from 64 kBaud to 4 MBaud in 1995. The increased speed and responsiveness, coupled with a rapidly increasing general awareness of the Internet and the WWW, resulted in an increased readiness to use the WWW for distributing teaching material. It was decided to conduct a much more extensive survey of Internet usage among SLAPA students. The results are presented here.

SLAPA Survey Results

These results come from a survey of student attitudes to electronic dissemination and gathering of information. The survey was conducted by questionnaire in May 1996. Eighty-three replies were received from students studying SLAPA. As part of their assessment, students were required to research topics such as on-line shopping and intellectual property issues associated with the Internet and to prepare presentations based on the results of their researches both orally and in the form of Web pages. The students taking the module were comprised of three groups: a group of 66 students who were nearly all final-year undergraduates; a group of 11 students who were finishing a 1-year postgraduate (MSc) conversion course; a group of 6 German exchange students from various courses. Most questions on the questionnaire were to be answered by circling a number in the range of 1 to 5. The suggested answers had a value or weight from most positive to most negative, which enabled weighted results to be calculated. The weights for the individual responses and the weighted totals are shown in the following tables. There was some evidence that some of the German students misunderstood the instructions, and the results for this group should be treated with caution. Not all students answered all questions, and so the numbers do not always add up.

The first set of questions asked about using the WWW to complete assignments. The results appear in Table 8.1

The weighted totals for the other groups were 3.30, 3.30, and 3.33 for the postgraduate conversion-course students and 3.17, 3.17, and 2.82 for the German students. These figures suggest that both groups were less happy with information

TABLE 8.1
Results: Use of WWW to Complete Assignments

Statement	*Agree Strongly (1)*	*Agree (2)*	*No Strong Opinion or Don't Know (3)*	*Disagree (4)*	*Disagree Strongly (5)*	*Weighted Total*
The information you found was useful.	17	28	14	6	1	2.18
The information was easy to find.	10	31	11	11	2	2.45
It was easy to print.	9	20	14	11	8	2.82

TABLE 8.2
Results: Questions About Library and CD-ROM

Group	*Had you used the following previously ?*						*Which would you prefer to use ?*		
	Library		CD-ROM		WWW		Library	CD-ROM	WWW
	Yes	No	Yes	No	Yes	No			
Undergrad	58	7	54	11	56	12	11	10	42
Postgrad	9	1	8	2	5	5	4	1	4
German	5	1	4	2	5	1	0	1	5

technology than were the final-year undergraduates, a conclusion supported by many of the following results; the postgraduate and German student groups were much smaller, however.

The two questions shown in Table 8.2 asked about previous experience using the library and CD-ROMs. A question on previous experience using the WWW was asked later in a separate part of the survey. The results have been included here for comparison purposes. Students were also asked which medium they preferred. The most striking figure in this set of results is the high preference shown for the use of the WWW to discover information. This finding is in conspicuous contrast to our earlier researches that had shown some preference for CD-ROMs; we concluded that these results were based on perceptions that the WWW was rather slow compared to CD-ROMs.

Questionnaire respondents were also asked whether they had any comments about the various media. Some of the comments are worth quoting. "WWW can be used from more locations" [22 respondents]: This response related to the fact that using conventional library resources clearly requires a physical trip to the library, and similar constraints apply to the use of CD-ROMs. CD-ROM-equipped personal computers are still uncommon at the university, and the university's stock of CD-ROMs is based in the library. Further, we do not have any networked CD-ROM drives so after a student has borrowed a CD-ROM, it is unavailable to anybody else until it is returned. It is likely that the falling price of CD-ROM-equipped computers will make them more common, but a library of CD-ROMs poses similar problems to a library of books. In contrast, the WWW is accessible from any of the university's computers, all of which are Internet connected, support some flavor of MS-Windows (Microsoft Corporation; http://www.microsoft.com/), and have either Netscape (Netscape Communications Corporation; http://www.netscape.com/) or Mosaic (National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign; http://www.ncsa.uiuc.edu/) WWW browsers as part of their standard software. Students with home computers can also access the WWW conveniently. The technology for multiple users to access a single CD-ROM does seem to be under active development.

"WWW is more likely to be up-to-date" [8 respondents]: This response might be more of a perception than a fact. At the time of the survey (May 1996), the WWW was still a rather new medium, and practically everything available on the

WWW would have been created in the 2 previous years. Whether this will continue to be the case remains to be seen; informal discussions with various colleagues have highlighted the WWW's impermanence as a potential problem for its long-term use. This problem is discussed later.

"WWW is often slow at peak times" [10 respondents]: This problem is familiar to all Internet users and is well discussed in the literature. Proposals and hints of proposals that future Internet users would have to pay more for reliable high-speed access are worrying to educators. Our particular context has several features that make this problem worse. For the SLAPA course, many identified resources were located in North America, hardly surprising as we share a common language (almost!). The transatlantic Internet links available to the U. K. academic community, however, are not particularly fast and are often congested. Further, in early mornings U. K. time, access is noticeably fast until American users start work (from about 11:00 AM U. K. time), when things start getting much slower and stay that way until about midnight U. K. time. Unfortunately, very few students choose to do their work during the faster access times.

TABLE 8.3
Results: Constructing and Using Web Pages

Statement	*Agree Strongly (1)*	*Agree (2)*	*No Strong Opinion or Don't Know (3)*	*Disagree (4)*	*Disagree Strongly (5)*	*Weighted Total*
Web pages were easy to construct.	17	28	12	7	1	2.18
Web pages were easy to put on the system.	18	18	14	12	1	2.37
The instructions supplied on the university's home pages were easy to follow.	9	25	14	5	2	2.38
I prefer producing a Web page to doing an oral presentation.	25	24	8	4	5	2.09
I prefer producing a Web page to writing an essay or report.	17	14	16	10	4	2.51
I learned more from looking at other students' Web pages than from listening to their presentations.	4	15	25	12	1	2.84

"WWW has more information than CD-ROM" [6 respondents]: This response was rather interesting. The students were given a set of links to identified resources, and they then followed the hyperlinks or in several cases went directly to the search engines. Any attempt to estimate the size of the WWW seems futile, but at the time of writing (November 1996), the larger search engines claimed to have indexed over 60 million pages with an average size of about 6 kilobytes, a total of 360 gigabytes of information. This is about 500 times the 600-megabyte maximum information on a CD-ROM, but many CD-ROMs include a much higher proportion of multimedia content than does the WWW. It is clearly unrealistic to expect a library of CD-ROMs to contain the same amount of information, yet questions about information duplication and information quality must concern teachers.

The next group of questions in the questionnaire covered the problems of constructing and using Web pages (see Table 8.3). The weighted totals for the other groups were 3.12, 3.62, 3.86, 2.75, 3.00, and 3.40 for the postgraduates and 3.00, 3.40, 3.00, 3.00, 3.33, and 3.00 for the German students. The most notable figures here are related to the difficulty encountered in creating and installing pages. The undergraduate group found this a fairly straightforward task, and thus these findings confirmed earlier results (e.g. Pitkow & Kehoe, 1996), but the postgraduates, with comparatively little experience, found it much more difficult (3.62 versus 2.37). The differences most probably arose because installing pages on "the system" required some knowledge of the Unix environment to which the postgraduates, who had only otherwise encountered MS-Windows-based software, were rather hostile.

The next three sets of questions were concerned with general information technology issues. The first set of questions dealt with attitudes to and use of E-mail (see Table 8.4). The weighted results for the other groups of students were 2.20, 2.55, 2.44, 2.44, and 3.38 for the postgraduates and 2.33, 1.40, 2.20, 2.20, and 4.00 for the German students. Several respondents expressed concern that assignments submitted by E-mail could go astray. It has certainly been our experience that misaddressed or misfiled E-mail is more difficult to recover than is paper that has gone astray, and any teacher contemplating receipt of assessments via E-mail should encourage students to keep backup copies on disk. The host operating system date stamps on such files should provide proof that the work was done before the handin date.

The next set of questions concerned the use of bulletin boards and the Usenet (see Table 8.5). The weighted totals for the other groups were 2.90, 2.60, 2.89, 3.11, 3.40, and 3.11 for the postgraduate students and 2.67, 2.33, 3.00, 2.33, 4.50, and 4.50 for the German students. Although a local Usenet newsgroup (wlv.scit.slapa) had been created as a forum for discussion of issues related to the SLAPA course, nobody posted anything to it. The reason is not clear although lack of publicity may well have been significant. The results in Table 8.5 show consistently positive attitudes about the use of bulletin boards, but it has proven difficult to get any class discussion going via this means. Other courses and

TABLE 8.4
Questions About Attitudes to E-mail

Statement	*Agree Strongly (1)*	*Agree (2)*	*No Strong Opinion or Don't Know (3)*	*Disagree (4)*	*Disagree Strongly (5)*	*Weighted Total*
I regularly read my E-mail.	41	13	1	4	6	1.78
I regularly use E-mail to send messages to lecturers.	22	14	14	12	3	2.38
Not all lecturers read their E-mail regularly.	11	18	25	5	3	2.53
E-mail is impersonal.	4	11	18	24	7	3.30
I think submitting assignments by E-mail is a good idea.	9	13	15	16	10	3.08

groups at the university have also attempted to set up group-specific discussion forums with a similar lack of success.

The next set of questions concerned general attitudes to and frequency of use of a variety of Internet-related tools and facilities (see Table 8.6). The weighted totals for the other groups were 4.00, 2.00, 2.56, 1.89, 1.78, 1.38, 1.78, 2.11, 3.44, and 3.22 for the postgraduate students and 3.83, 4.17, 3.17, 3.00, 2.00, 2.00, 2.83, 2.00, 4.00, and 3.83 for the German students. The results shown here clearly indicate the extent to which the WWW has already displaced earlier Internet resource discovery and access tools such as anonymous FTP and Gopher in the students' perception of how to use the Internet. This cohort of students had not encountered the use of the Internet in any other courses; later cohorts, however, will have encountered the Internet in a course called Business Data Communications, in which they are required to send each other E-mail and also to use the WWW to discover resources relevant to the course.

The Way Forward – Resources

Our experiences have suggested that students like to use the Internet and more especially, the WWW. Individual teachers and the educational community must however address several significant problems before the Internet's potential to enhance the educational experience can be fully unlocked. It is far from clear that unrestricted exposure to the multiplicity of resources available is an effective and efficient use of student time. Most teachers think that it is not, and we have often found that a number of students have wandered off into some odd corners of the WWW. There seems to be no way to stop this pattern of activity, although the presentation of a carefully selected set of resources can be very helpful.

Unfortunately, teachers must prepare the selection, which is probably presented as a set of links on a course-related page, and must spend considerable time and effort eliminating irrelevant and out-of-date items. A teacher approaching this task probably starts by using one of the major WWW search engines, but effective use of these is not simple, and no useful published guides help to use such services effectively.

The design of most search engines seems to owe little or nothing to library science and the skills acquired by librarians in organizing and cataloguing information. (See, for example Mannoni, 1996.) Unless students have plenty of time or have had considerable tutoring in using search engines effectively, they cannot be told to find out about something "on the Internet." Such an activity is likely to be frustrating and time wasting. Some search engines offer directories or hierarchies of directories; this approach is much closer to traditional classified library catalogues, but the selection of categories is sometimes rather arbitrary. In

TABLE 8.5
Questions About Bulletin Boards and Usenet

Statement	*Agree Strongly (1)*	*Agree (2)*	*No Strong Opinion or Don't Know (3)*	*Disagree (4)*	*Disagree Strongly (5)*	*Weighted Total*
Bulletin boards and Usenet newsgroups provide a good means for people to ask for help with problems.	11	31	15	5	2	2.31
Only the most extroverted students and staff post messages to bulletin boards and newsgroups.	4	22	16	15	7	2.98
With a bulletin board everyone has a say, not just the most vociferous students.	9	31	15	8	1	2.39
I would like to take part in SLAPA class discussions on the SLAPA bulletin board.	7	16	28	10	3	2.78
It is a good idea to post marks and other messages on the class bulletin board.	12	19	15	12	6	2.69
I find it difficult to use a bulletin board.	4	10	17	16	15	3.45

TABLE 8.6
Questions About Internet Tools and Facilities

Statement	*Never (1)*	*Rarely. Tried It Once (2)*	*Occasionally. Less Than Once a Week (3)*	*Frequently. At Least Once a Week (4)*	*Daily (5)*	*Weighted Total*
I use electronic mail.	1	0	10	13	41	4.43
I use Telnet to login to remote computers.	7	5	15	17	21	3.62
I use FTP to transfer files to and from other computers.	11	6	20	16	10	3.13
I use anonymous FTP to obtain data and resources on public archives.	20	10	16	11	6	2.57
I use Archie to locate resources on public archives.	33	15	7	5	3	1.89
I subscribe to one or more mailing lists.	30	12	10	3	8	2.16
I read one or more Usenet newsgroups.	18	12	19	5	12	2.71
I use Gopher to access remote resources.	31	16	9	4	3	1.92
I use WWW to access remote resources.	3	4	12	24	17	3.80
I use WWW search engines to locate remote resources.	6	4	14	20	20	3.69

an institution of any reasonable size, it is probably helpful to identify an "expert" in the use of search engines. Such a person may well be a librarian or a person with training in librarianship. Alternatively, a teacher may choose to refer to an increasing number of Internet and WWW resource lists being prepared and published by educators. Even if these match the course structure and content, many doubts raised in the following paragraphs must apply to such lists.

Internet-based resources remain the property of their authors who can modify and remove them at any time, without previous notice. In marked contrast, library copies of resources remain in the library until removed by librarians. Once a resource has been transferred to a library, an author no longer has any control

over it. There has been a little discussion about preparing a proper archive of the WWW, but as far as we know, there are no definite proposals.

There are also significant unresolved problems with intellectual property rights associated with material "published" on the WWW. The ease with which students can cut and paste material from their Web browsers into their own word-processed assignments is also a cause for concern for teachers who may well inadvertently reward, as original, material that has been copied from an unfamiliar source. An important point for teachers to consider if they intend to require students to create Web pages as part of an assessment is the problem of plagiarism and the possibility of changes being made after handin dates. Once a Web page has been created, it is, in general, public and can be viewed and copied by anybody. Students who begin to work later than others can simply copy other people's already published results. We considered using special software to make automatic copies of pages for assessment purposes, but the software was not entirely successful. Another possibility would be to require students to submit the text of their pages via E-mail or on a floppy disk.

It is often difficult to judge the quality of a resource. With conventional published material, the reputation of the publisher and the mechanism of reviewing provide a reasonable guarantee of quality. There is no such mechanism on the WWW, and it is not always easy to tell whether a page is a result of peer-group reviewed research or wild speculation perhaps supported by groups with extreme views.

There is an increasing tendency to use the WWW to publish material in a bewildering variety of formats, particularly under the banner of multimedia. This practice may be acceptable in a school or college context where it is possible to ensure that all machines are equipped with all the relevant viewers, but it can cause severe problems for home-based students who may not have the latest, fastest, most powerful machines.

The Way Forward – Interactive Use

Since the WWW was first created, it has had the ability to support interaction via hyperlinks. More recent versions have included the forms mechanism and the associated common gateway interface (CGI) mechanism that allows information gathered from remote browsers to be processed by local programs on the server computer, and thus provides a basis for client-server computing without the need to distribute special client programs. So-called sensitive maps provide another interactive mechanism. All these tools and methods seem to have been little used in education, possibly because of the significant computing skills needed.

The older and computationally simpler technique of multiple choice questions answered via selecting a hyperlink to one of several possible answers is also very rarely seen on the WWW. The problem here seems to be partly related to the difficulty of devising such tests, although some interesting automated approaches have been reported (Krile, Juells, & Vetter, 1996).

Web Resources

Parts of this document and links to related material can be found on the WWW at http://www.scit.wlv.ac.uk/results/.

REFERENCES

Burden, P., & Davies, J. (1996). Appropriateness of the Internet as a teaching resource. In *Proceedings of leading edge teaching technologies - LETT '96* (pp. 24–27). Victoria, BC: University of Victoria.

Dern, D. P. (1994). *The Internet guide for new users*. New York: McGraw-Hill.

Krile, T., Juell, P., & Vetter, R. (1996). Pioneering on the new frontiers of education. *Computer, 29*(5), 112–114.

Mannoni, B. (1996). Bringing museums online. *Communications of the Association for Computing Machinery , 39*(6), 100–105.

Pitkow, J. E., & Kehoe, C. M. (1996). Emerging trends in the WWW user population. *Communications of the Association for Computing Machinery , 39*(6), 106–108.

9 Easing the Learning Curve for Novice Web Users

Julie Ratner
The Mathworks, Inc.

There are three reasons that research on college students' attitudes toward learning on the World Wide Web (WWW) is important. Instructional and user interface designers (UI) need data on novice Web users, particularly learners in the domain of education, to build more user-friendly browser interfaces. Web browser designers who are creating new features and enhancing existing features for online services also need this research to understand how to make browsers more intuitive. Novices' attitudes toward computer learning affect their performance; as textbooks are replaced by computers as teaching delivery mechanisms, authors and publishers must consider how attitudes affect comprehension.

In this chapter, I profile novice Web users and describe characteristics unique to this target population, characteristics based on their attitudes and their performance when completing typical scenarios using the Netscape (Netscape Communications Corporation www.netscape.com) browser. Because of their inexperience and overconfidence, novices accessing curriculum via the WWW face impediments in using a browser efficiently and show decreased levels of comprehension. A list of implications offers ways to ease the learning curve for novice Web users exposed to Web technology at college levels.

To pinpoint attitudinal differences between novices and expert Web users, I (Ratner, 1996) conducted a study that analyzed three scores: usability, attitude, and comprehension. Two scores, browser usability and Web attitude, are discussed in this chapter. Participants were considered novices if they had never used the WWW; they were considered experienced if they had used the WWW at least two times. The usability score measured how many scenarios testers successfully completed and rated whether navigation and control over the display were easy and intuitive. The attitude score evaluated the testers' receptivity to incorporating Web technology into their university programs.

I used the Web Attitude Measure (WAM) to determine how receptive college learners were to using the WWW for curriculum delivery and comprehension testing. The hour-long session consisted of a sequence of tasks: signing a consent form, completing five scenarios using the browser, receiving instruction on the five scenarios, completing the WAM, reading curriculum on the WWW, closing the network connection, and finally, answering questions on the

curriculum. No tasks had time limits, to encourage novice users to experiment and explore the browser functionality without pressure. The results indicated that previous experience using the Netscape browser significantly increased usability, attitude, and comprehension.

Novice Web User Demographics

Ninety-seven University of New Mexico (UNM) undergraduate and postgraduate students participated in the study; novice Web users predominated because I had advertised for people with no or little Web experience. Only 32 of the 65 subjects had ever used the WWW before the research session. The baseline requirements to participate were a minimum age of 17 years, enrollment at UNM, computer keyboard skills, and 1 hour for testing. Subjects filled out a questionnaire with demographic data to distinguish those with experience in computers and E-mail from those with Web experience. There was a nearly even distribution by gender and age, with 52 subjects under 25 and 41 over, as indicated in Table 9.1. The amount of computer and E-mail experience among novice Web users confirmed that novice Web users are computer and E-mail literate and that distinct skills are required for Web literacy.

To establish usage patterns, experts were asked specific questions about the amounts of time they spent on the WWW weekly and what Web sites they visited. Out of 32 experts, nineteen previously used the Netscape browser, fourteen logged on from school, and sixteen browsed articles, references, papers, and weather sites. Many respondents commented that they "surf" and "browse" for "entertainment" and "crafts" as they looked for Web pages relating to their hobbies or recreational interests.

Novice Web Users Speak Out on Browser Usability

To evaluate the performance usability of the Netscape browser (version 1.1), I asked participants to perform certain tasks considered typical for the WWW. The test session started with the browser open to the UNM home page. Participants then tested five different browser interface features that followed a logical sequence from simple to more difficult. After each scenario, participants

TABLE 9.1
Demographic Analysis of Sample

	Gender	*Computer Usage*	*E-mail Usage*	*Avg Hours Browse Web (Per Session)*
Novice	M 21	Under 1 yr 15	Daily 7	
	F 44	1 - 3 yrs 25	Weekly 30	
		4 - 6 yrs 15	Monthly 3	
		7 + yrs 10	Rarely 25	
Experienced	M 14	Under 1 yr 4	Daily 13	Under 1 hr 20
	F 18	1 - 3 yrs 7	Weekly 15	1 - 2 hrs 9
		4 - 6 yrs 8	Monthly 0	3 + hrs 3
		7 + yrs 13	Rarely 4	

TABLE 9.2
Actual Usability Scenarios Correct (%) by Novice and Experienced Users

Usability Scenarios	Novice (67%)	Exp (33%)
Scenario 1: Go to the Libraries page.	75%	88%
Scenario 2: Return to the previous page.	57%	84%
Scenario 3: Go to the University of New Mexico Web page.	63%	84%
What is the Web address?	57%	88%
Scenario 4: Change the display font to Large.	11%	38%
Scenario 5: Look for Web sites relating to Psychology.	37%	60%

wrote down how they completed the task and then chose yes or no to answer whether they found the task easy.

Among all the testers, the *perception* of usability was very high $F(2,93) = 12.1349$, $p < .000$, but the actual *performance* was low. This dichotomy between user perception and performance clearly illustrates that users are unaware of their usability problems with the browser interface. Novice and expert users had nearly identical mean perceived usability scores; the average score was 4.4 out of a total possible score of 5. Actual usability scores revealed this discrepancy between perceived usability and actual performance, as indicated in the responses to the five scenarios.

Testers were first asked to connect to a Web page, then to return to the previous page, then to find the UNM home page address or uniform resource locator (URL), then to change the font size, and finally to complete a search for other Web sites relating to education or psychology. Novices performed poorly on these tasks because of the usability of the browser (see Table 9.2). The percentage of successfully completed scenarios was consistently lower for novice users compared to experienced users. Dramatically low completion rates were tabulated even for experienced users for the two more difficult scenarios. These low usability scores indicate that both advanced navigation and browser display control features cause confusion for users. User questions and comments were collected during and after the scenarios by using a checklist to record behavior patterns in relation to browser usability. Preference settings to change the display fonts are one instance of browser design incongruous with users' cognitive models.

All participants had trouble changing the font size and finding preferences off the main menu bar. Both novice and expert users erroneously looked for menu bar labels consistent with word-processing software categories, like Edit or Format, to change the browser text font. This fact was indicated by user comments recorded during the usability trials: "I'm looking for the format option to change the font size," "Where is the Edit choice to get to fonts?" The general preferences menu category and font option were illogically grouped with other types of browser application settings and hidden three levels deep. Novices and experts could not complete the task because they did not make the distinction between the page display and the browser settings. It was clear from the comments and questions that browser usability was not intuitive for other reasons

as well. Even when testers successfully completed the scenario, it was not obvious to them why the text font size changed and the image map font did not; users did not understand that font in an image map is part of the graphic.

There was also confusion between local network and Internet connections. On the basis of responses to the last scenario, there was conclusive evidence that participants did not make a distinction between a search on the local university Web page and one on the Web. Nearly all the users who thought they had completed an Internet search, had actually written a Web address local to the university as a response. Few of the users, novice or expert, actually found the browser search icon or any of the Internet search engines.

While the participants were completing the scenarios, I observed them, and completed a checklist of their behaviors and the mechanical problems that might have affected their completion rate. I also recorded participants' specific questions for analysis. Most problems observed were mechanical: The mouse failed, the Internet connection broke, an Insufficient Memory message appeared. Typically, when these problems arose, the person was moved to another machine to continue the trials without delay. Non-mechanical difficulties were due to students' unfamiliarity with the interface. At every interaction, subjects were reminded that the purpose of the scenarios was to test the performance usability of the browser, Netscape; if they could not complete a trial, they were to write down what they had tried and then move on to the next trial. I continually reassured the testers that any difficulties in completing the scenarios meant that the browser design needed to be refined and improved to be more intuitive. Despite my message, many testers wasted time on the third and fourth scenarios and never tried the fifth and hardest scenario.

The most typical questions students asked pointed to the key problems with the interface and the WWW. These included:

- How do I go to the main menu?
- How do I find fonts? I'm looking for Options on the menu.
- What do I do if I can't answer questions?
- Is Netscape considered a browser?
- What do you mean by *navigate*? Am I on the last page? What happens when I get stuck on a screen and can't get back?
- I can't find the home page. Is this screen considered a Web page? I don't know where to find a Web page.
- (*Closes Netscape*) "I'm not sure how I did that."
- (*Clicks icon*) "Is it that simple?"
- (*Clicks bullet, not hyperlink*) How do I go to the next page?

Even after instruction on how to complete all the scenarios, novice users found it difficult to access the curriculum. The destination address at the bottom corner of the browser window was too subtle a cue that different links navigated to other local or Internet pages. The jargon used in the scenario questions and on the

browser interface, such as browser, Web, home page, hyperlink, and Web address, impeded usability and prevented users from focusing on the task at hand.

Novice Web User Issues for Interface Design

Novice users rarely distinguish between the browser interface and the document; therefore, they are usually perplexed when the two behave differently (Nabkel & Shafir, 1995). Nabkel and Shafir explained that "the boundary between application interfaces and their interactive content has blurred and continues to disintegrate." Unlike word-processing or desktop-publishing software, browser interfaces do not allow designers to control font style and size, alignment, window size, and other aspects of document design (December & Randall, 1994).

The fact that browser interface design affects users' cognition, behavior, and attitude presents many challenges to browser interface designers. The overwhelming interest in the Internet and the multi-platform environment in which it exists generate a wide variety of screen design and display problems. Monitor screens lack any spatial cues, and computer-based instruction has no physical properties, like page numbers in a paper textbook (Miller, 1985). Without page numbers and chapter layout divisions to assist in dividing information, informational grouping on screen does not fall into discrete display units. Visually, text, graphics, and the way the browser interface window displays them are intricately interwoven. Continued research in human factors would help meet these design challenges and facilitate novice Web users' learning via the WWW. Composing information in a browser interface also poses complex problems affecting novice user attitudes.

Novice Web User Attitudes

The Web Attitude Measure (WAM) is an instrument based on Kay's (1989) Computer Attitude Measure (CAM), that was modified to focus specifically on the WWW, rather than on computers in general. The WAM is divided into four sub-scales: behavior, affect, attitude, adjective. The relation between usability and the attitude sub-scale illuminates interesting points about novice Web users. When

TABLE 9.3
Standard Deviation and Effect on Affective Scores for Experienced Users

Affective Questions	*SD*	*F*
1. I feel I need an experienced person nearby when I use the WWW.	1.909	.000
2. I need someone to tell me the best way to use the WWW.	1.738	.000
3. I could probably teach myself most of the things I need to know about the WWW.	1.846	.000
4. I can make the WWW do what I want it to do.	1.642	.000
5. I am in complete control when I use the WWW.	1.734	.000
6. If I had a problem using the WWW, I could solve it one way or another.	1.596	.042
7. I would prefer to learn new Web browser software packages on my own.	1.644	.030

a one-way Analysis of Variance was performed on affective scores by web experience, there is a significant difference by experience $F(2,94)=14.9622$, $p<$ .000. Every question in the affective sub-scale, listed in Table 9.3, confirms a significant effect on affective scores for experienced users.

Novice users wanted an experienced user nearby to answer questions, someone to show them how to use the Web, and to learn new browser packages. They did not feel in control or confident using the interface without assistance, despite their years of experience using computers and E-mail. They did not feel confident using the WWW despite their perceived satisfaction with the browser during the usability scenarios. Novice users made a distinction between the WWW and the browser experience, an indication that their uncertainty is focused on the mechanics of connecting to the WWW and of self-teaching to use Web software. The low affective scores among novice users confirmed that interfaces must be better designed for inexperienced users and instruction on how to use the WWW must be provided before using the Internet to deliver curriculum in classrooms.

Novice Web Users May Also Be Disadvantaged Learners

There are many ways to approach the issue of usability, the most common of which is to focus on information accessibility in the context of online documentation. Nabkel & Shafir (1995) used the phrase "metaphor and interaction dissonance" to describe the challenge of viewing dynamic, interactive movable objects (documents) with a static browser application. Developers of online curriculum can build on the strategies typically used for reading static, linear texts by providing Web users with spatial cues. These strategies include skimming, look backs, and rereading text to create a cognitive schematic of the information, with chapter headings and graphics as landmarks (Waller, 1986). The usability of an online document is directly related to these cohesion cues and ultimately to reading comprehension (Halliday & Hasan, 1976).

In online curriculum design, tasks, not the tools, are central (Norman, 1990). Designers should not focus on the interface, but on the ease with which users can access informational concepts. Yet information cannot be separated from the interface (Nabkel & Shafir, 1995). If a user has trouble accessing the information or performing a task, then both the interface and the document structure are to blame, and both will have a negative effect on the user's attitude.

SUMMARY

Many students are first introduced to the WWW in the classroom. Researchers have agreed about the benefits of incorporating Web technology in the classroom to expand resources and to allow worldwide communication. Technology, however must be incorporated with novice users in mind. To prevent low comprehension and further division of hardware haves and have nots, faculty must consider the implication of merely transferring paper-based materials to the

WWW. Following is a list of implications for novice Web users in any classroom environment:

1. Novice Web users need to be provided instruction on how to use a browser and how to install helper applications and plug ins.
2. Novice Web users need to be shown how the client - server architecture operates and how to recognize and resolve connection problems and error messages.
3. Faculty need to poll classes for novice and expert Web users, and pair novices and experts for the first few classroom projects in order that inexperienced users can gain confidence with Web technology.
4. Faculty need to clarify different connection trade offs from home and school and various browser features with text-only capabilities for faster downloading options.
5. Curriculum designers must include file sizes, particularly for large images and, audio and video clips, so that students can make informed decisions before downloading files.
6. Faculty are encouraged to poll classes for learning styles to be aware of positive and negative attitudes toward the WWW as a delivery vehicle.
7. Faculty need to poll classes for hardware and software access; they must be familiar with university computing center hours to ensure that all students have sufficient access to classroom materials presented only on the WWW.
8. Faculty need to integrate classroom goals and direct Web-based projects to ensure that students' time spent on the Internet is focused on browsing instead of on recreation.
9. Faculty need to redesign curriculum for Web delivery; they must also increasingly integrate images with captions and associate Web links to maximize the true potential of this technology.
10. Faculty should anticipate anxiety among novice Web users in classes using Web technology.
11. Faculty should anticipate over confidence among novice Web users in classes using Web technology.
12. Curriculum designers need to accommodate different browser limitations and configuration problems for students to ensure that all students can access the instruction.
13. Curriculum designers need to provide navigation in addition to browser navigation to facilitate associative learning across sections and within sections of instruction.

Novice and experienced Web users have provided critical data about the usability problems of interfaces. Most neophytes on the WWW took about 5 hours to learn from friends or in a course how to use a browser and author an Hypertext Markup Language (HTML) Web page (Pitkow & Kehoe, 1996). By observing

and surveying this novice population, collected data help identify how to shorten the 5-hour learning curve and how to best meet instructional objectives on the WWW.

To successfully incorporate the WWW in education, designers and educational institutions must take human factors research into account. Because overly complex browser interface features detract from comprehension, browser designers must simplify the layout, graphic file format and size, and document structure when developing curriculum for novice Web users. Faculty must provide an introduction to the WWW before delivering information using this technology.

Netscape (version 1.1) was difficult for novice users to navigate: Subjects could not complete the five browser scenarios successfully. Users' familiarity with Web interface and ease of navigation using the browser increased usability performance and attitude. As demonstrated by observed behaviors and poor performance on usability trials, novice users' difficulty in accessing curriculum impeded comprehension. Thus, the browser, by impeding users' access to the curriculum, negatively affected comprehension of the material as well. Because the newest version of Netscape (now 4.0) is even more complex than the version (1.1) tested during these trials, more research on the usability of the browser by novice users is needed. As the Web user population continues to increase for all age groups, interface usability testing must continue and browser redesign take place to accommodate broader, more diverse audiences. Once browser interfaces are designed for non technical and younger populations, the WWW will become an even better learning resource.

REFERENCES

December, J., & Randall, N. (1994). *The world wide web unleashed.* IN: Sams.

Halliday, M. A. K., & Hasan, R. (1976). *Cohesion in English*. London: Longman.

Kay, R. H. (1989, Summer). A practical and theoretical approach to assessing computer attitudes: The computer attitude measure. *Journal of Research on Computing in Education*, 456–463.

Miller, R. (1985). "?," *Visible Language*, *XIX*(2), 144–150.

Nabkel, J., & Shafir, E. (1995). Blazing the trail: Design considerations for interactive information pioneers. *Special Interest Group Computer Human Interface (SIGCHI) Bulletin*, *27*(1), 45–54.

Norman, D. A. (1990). Why interfaces don't work. In B. Laurel (Ed.), *The art of human - computer interface design* (pp. 209–220). New York: Addison-Wesley.

Pitkow, J., & Kehoe, C. (1996). *GVU's 4th world-wide web user survey*, Georgia Institute of Technology, http://www.cc.gatech.edu/gvu/user_surveys/survey-10-1995/.

Ratner, J. A. (1996). *Features influencing the use of the world wide web for learning: Usability, attitudes, layout design*. Unpublished dissertation, University of New Mexico.

Waller, R. (1986). What electronic books do better than. *Information Design Journal*, *5*, 72–75. Harper Collins.

Part III: WEB DESIGN GUIDELINES AND DEVELOPMENT PROCESSES

10 Using Web and Traditional Style Guides to Design Web Interfaces

Eric Grose
Chris Forsythe
Sandia National Laboratories

Julie Ratner
The MathWorks, Inc.

We would have preferred to say that Sandia National Laboratories had fully anticipated the explosive growth of Web technologies. Instead, software projects in the midst of implementation have been completely scrapped because they could not compete with the solutions introduced by Web interfaces. That Web technologies took Sandia somewhat by surprise is no source of shame, however; the same fate befell many other companies, and in many respects, Sandia has been a leader in the development of Web-based solutions.

Like most large organizations, Sandia had struggled to deliver corporate software standards in a heterogeneous computing environment that includes many Unix, personal computer, and Macintosh users. Platform independence had motivated purchases of products to make software portable to different platforms, but this portability usually came at the expense of the user interface (UI). When Web technologies appeared as a cross-platform software solution, Sandia was quick to adopt this paradigm.

Early in 1994, Sandia launched an "intranet" to support a re-engineering of corporate processes that would enhance the overall productivity of the Labs. Over the past couple of years, corporate applications—vacation balances, travel forms, reimbursement vouchers and training classes, to name a few—have been made available to every employee through the intranet. In the past, these transactions would have been available to Sandians only through a human intermediary and nearly always on paper. As a result, the intranet has transformed many aspects of Sandia's corporate culture.

This proliferation of web applications has changed the environment for Sandia's Human Factors Department, with increasing levels of effort committed to Web interfaces. Working with World Wide Web (WWW) interfaces has been challenging for several reasons. For one, the browsers used to access the WWW are limited in their ability to deliver a superior interface. Consequently, as usability specialists, we have had to occasionally grimace at the clumsy interfaces that result. But, Web capabilities are rapidly improving (frames, JAVA) and these

improvements have introduced challenges of their own. It is difficult to derive definitive interface solutions when the technology may substantially change in 1 year.

We believe that computers, and the WWW in particular, have in many ways failed to meet their potential to enable users to work smarter and more productively. This chapter describes research and development addressing Web style guides, an area that we have pursued in our efforts to improve decisions being made about interface designs for Sandia's intranet.

Since the graphic desktop computer revolution in the 1980s, our bookshelves have filled with documents of recommendations about how to design computer screens. When the WWW appeared, we and many others were unsure how traditional UI design guidelines related to this new and exciting software tool. Much of the guidance was obviously valid and some of the initial multicolored screens reminded us of the days when software developers were first given access to thousands of colors. But the realities of a hyperlinked world—client–server architectures and local versus remote access, to name a few—negated some of our knowledge.

As the WWW has grown, there has been a proliferation of how-to documents on the WWW itself. Most are small and narrow and are written for a specific audience. We viewed these Web style guides with some curiosity. Did they contain good UI information? Should we recommend or discourage their use? To answer these questions, we conducted a series of studies that addressed different facets of Web style guides.

DEVELOPMENT OF STYLE GUIDES: WEB VERSUS TRADITIONAL

Our first study sought to identify who was writing these Web style guides and why. We sought to answer questions about the sponsorship and publishing environments that promoted these documents.

Ten authors of style guides, five Web and five traditional, responded to a survey that asked them about their educational and professional backgrounds. Authors were asked to describe why they created a style guide, who was the intended audience, and who sponsored its development. Questions were also asked about the time it took to create the style guide, the resources used, and how often it was updated.

Survey results from Web and traditional style guide authors revealed distinct differences. All Web style guide authors were from educational environments; they described the development process as informal, with an average of 3.26 weeks spent writing the style guide and with revisions prompted by feedback from users or impending conference submissions. In contrast, the traditional style guide authors were from military, government, or corporate organizations. Three of five specified that a formal process was employed to develop the style guide, with an average of 54.2 weeks for the initial version and with subsequent versions following annual reviews with structured working groups.

Web style guide authors explained that resources to write the guide drew on the fields of Human Factors, marketing, graphic arts, and public relations. Traditional guidelines were based primarily on Human Factors, cognitive psychology, and military sources. The audience for traditional style guides was defined as programmers, interface designers, and technical experts. The audience for Web style guides seemed less clear: In only one instance did an author identify a specific audience for the document — the staff at the university the author was associated with.

This initial study indicated substantial differences between Web and traditional style guides and prompted further analysis.

STYLE GUIDE CONTENT

The second study compared the content of Web and traditional style guides to examine the extent that Web style recommendations included the most essential of the traditional ones, as judged by a group of Human Factors practitioners.

Style guides were sampled in the following manner. First, 21 Web sites offering Web interface recommendations were identified in fall 1995. From these sources, a set of 357 unique recommendations was obtained. Table 10.1 lists the most frequently cited recommendations from the Web style guides. Second, recommendations from traditional style guides were collected from five printed sources. Style recommendations not applicable to the WWW or specific to a particular platform or operating system were discarded. This process resulted in a collection of 270 unique interface recommendations.

Style recommendations from both Web and traditional style guides were assigned to 1 of 20 categories and compared. Figure 10.1 shows the distribution of recommendations across categories. Four categories accounted for 63% of the Web recommendations: Common Look and Feel (CL&F), Information Display, Navigation, and Labels.

TABLE 10.1
Most Frequently Cited Recommendations From Web Style Guides

Recommendations	*# Style Guides*
Signature should be included at bottom of Web page.	12
"Here" or similar words should not be used to designate a link.	12
Pages should be time stamped and dated.	11
Every page should have a title.	11
Design for all platforms and browsers.	10
Menu bar or buttons should be provided for navigation.	8
Text used for links should be descriptive and meaningful.	7
Pages should have common look and feel.	7

As shown in FIG. 10.1, the categorization of recommendations reveals some overlap but also considerable differences between Web and traditional style guides. A chi-square test of independence showed a significant difference in the distributions of recommendations across categories ($X^2 = 245.5$, $p < 0.05$; $df = 19$). The biggest differences arose from the emphasis on information display and consideration of help, data entry, and message boxes in traditional style guides.

Another comparison assessed the degree to which traditional style recommendations were addressed in Web style guides. Of the 270 Web-relevant recommendations found in traditional style guides, only 53 (20%) also appeared in Web style guides, with another 20 partially addressed in the Web style guides. The greatest overlap occurred for recommendations categorized as Common Look and Feel (15 of 16), Menu Design (11 of 17), and Labels (15 of 25).

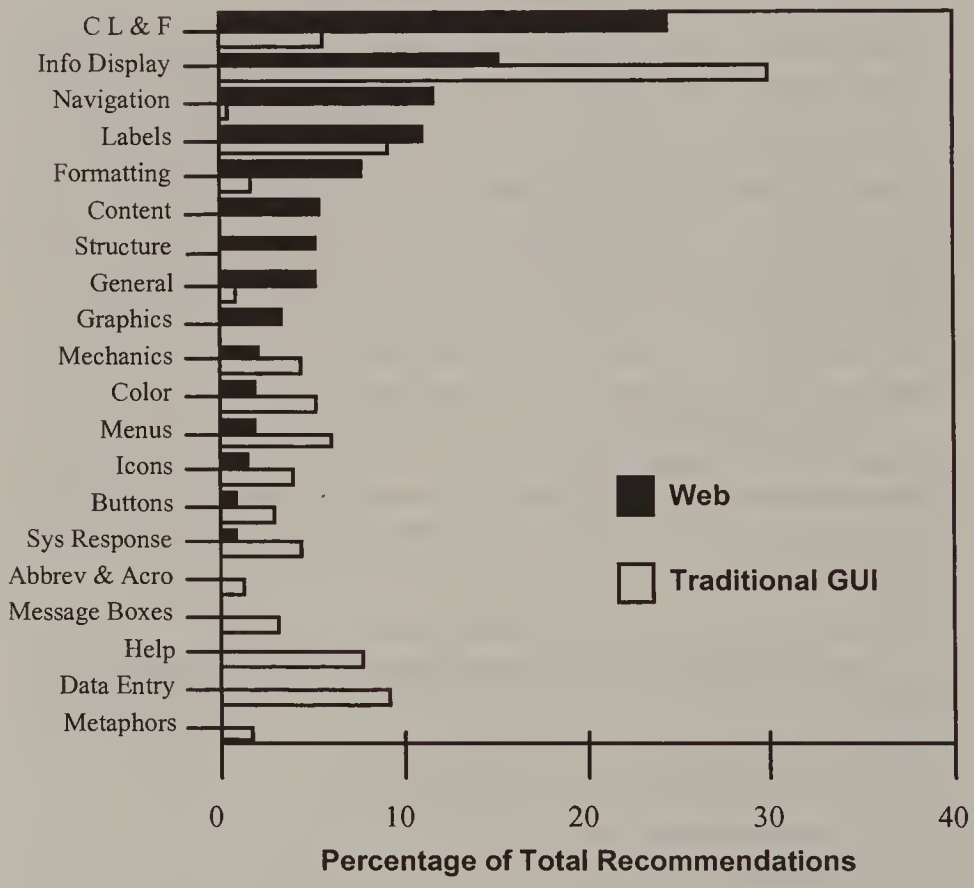


FIG. 10.1. Distribution of style guide recommendations across categories.

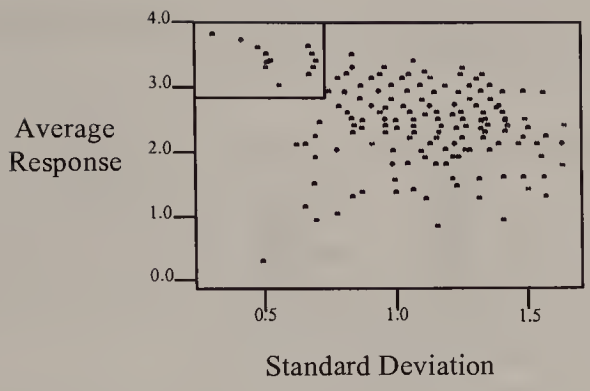


FIG. 10.2. Plot of the mean and standard deviation of "essentialness" ratings for each of 270 recommendations from traditional user interface guidelines.

ASSESSMENT OF RECOMMENDATIONS ESSENTIAL TO USABILITY

To assess the representation of the "most essential" traditional style recommendations in the Web style recommendations, it was first necessary to determine which recommendations were most essential. Eleven practicing Human Factors professionals completed a survey asking them to review the 270 recommendations obtained from the traditional style guides. For each recommendation, participants gave their opinion as to how essential they thought the recommendation was to the usability of a software UI. Answers were indicated by circling an integer on a scale from 0 (nonessential) to 4 (absolutely essential). Here is a sample question:

Labels should be placed in close proximity to their related data fields.

0	1	2	3	4
Nonessential		Important		Essential

Descriptive statistics were used to identify items that participants generally agreed were essential to usability. These essential recommendations from traditional sources were compared to the Web recommendations to determine how well they were accounted for in the Web domain.

A plot of the data obtained from the surveys revealed 27 recommendations whose mean score was 3.0 or greater and whose standard deviation was less than 0.8. These items received consistently high ratings from the experts. Figure 10.2. provides a plot of the survey data by mean and standard deviation. The boxed area reflects the most essential guidelines. Of these most essential recommendations, only one—Push buttons related to overall window functionality should be located at bottom of screen—was found in any of the Web style guides. Table 10.2 shows

the ten highest ranked recommendations from traditional style guides with the complete list of 27 appearing in Appendix 10.A.

EVALUATION OF GUIDELINES BASED ON ACTUAL FINDINGS

As a further assessment of their practical effectiveness, Web style recommendations were compared to actual Human Factors reviews of the Web-based corporate information infrastructure at Sandia National Laboratories. (The reviewers were knowledgeable of the traditional style guides, but all reviews occurred prior to exposure to any Web style guides.)

There are generally two different Web interfaces. First are interfaces that derive information for users by performing routines (e.g., querying databases) or allow submission of information (e.g., online forms); this chapter refers to such interfaces as applications. Examples of applications are an electronic phone book or a conference room scheduler. The second Web interface is relatively static and displays information or serves as a waypoint for navigating to other Web pages; examples are a corporate newsletter, engineering procedures, or departmental home pages. Findings from these two Web interfaces were compared to Web style guides to assess how well the guides accounted for deficiencies identified by the Human Factors specialists.

TABLE 10.2

Ten Recommendations Judged Most Essential by Human Factors Experts With Mean Ratings of Essentialness and Standard Deviations (SD)

Recommendation	*Mean*	*SD*
Data entry should be accomplished through a distinct action.	3.7	0.49
Feedback should be provided about log-on procedures to indicate log-on status.	3.6	0.67
Each color should represent only one category of displayed information.	3.5	0.52
Displayed information should be readable to the degree of accuracy required by the task.	3.5	0.79
Push buttons related to overall window functionality should be located at bottom of screen.	3.5	0.79
Axis of line graphs should be clearly labeled with units of measurement, if appropriate.	3.4	0.51
Where a user has the option of deleting an entry that is not displayed, the user should have the option of viewing the entry before confirming deletion.	3.4	0.52
Information should be presented in a directly usable format that does not require decoding, interpretation, or calculation.	3.4	0.67
Text should be readable from normal viewing distances.	3.4	0.67
When data have a naturally occurring order, it should be reflected in the organization of data fields.	3.4	0.67

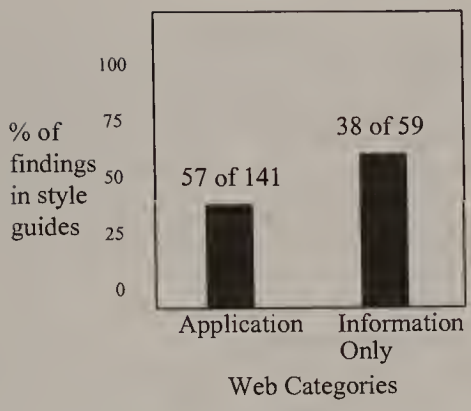


FIG. 10.3. Percentage of findings accounted for in web style guides.

In the review of corporate Web-based applications, only 57 of 141 findings (40%) corresponded to recommendations presented in the Web style guides. In contrast, for Web pages primarily with information content, 38 of 59 findings (64%) were addressed in Web style guides (see FIG. 10.3). These proportions are statistically different at the $p < 0.001$ level of significance.

PRACTICAL USE OF WEB INTERFACE DESIGN GUIDELINES

In total, several hundred people at Sandia have, are currently, or will eventually develop materials to be placed on the corporate intranet. The Human Factors Department has worked diligently to influence the design of these Web-based products and has pursued various approaches to improve the usability of Web products. We have written reports and taught Human Factors courses, but mostly we have simply reviewed and tested interfaces. Because we are a very small group (5, only 3 of whom are dedicated to working with Web applications), we have had to find other ways to impart Human Factors guidance to Web developers. One of these is publishing Web UI design guidelines for use by the Web development community.

Software developers are busy, practical people who look for ways to make their complex jobs easier. Those who have forayed into the world of Human Factors have often been intimidated by the volume of guidance. If they have no other use for Human Factors practitioners, they certainly want us to interpret this glut of information, to give them something they can comprehend and use.

In spring 1996, a set of guidelines was derived from the work previously described. Guidelines from the Web and traditional style guides were combined, with a total of 574 unique recommendations. Each of these 574 recommendations was scored on the following criteria:

> **Practicality**: The labor required to implement the recommendation should be in proportion to the benefit anticipated from it.

Verifiability: The developer should be able to tell that the recommendation has been implemented.
Recognition: Multiple sources should include the recommendation.
Criticality: The recommendation should strongly effect usability.
Relevance: The recommendation should be relevant for many Web pages.
Occurrence: On the basis of evaluations of Web pages, the recommendation should have been observed to have been violated frequently.
Clarity: No special knowledge or training should be necessary to implement the recommendation.
Constraint: The recommendation should not unduly constrain the ability of Web page designers to seek innovative solutions.

On the basis of these criteria, two sets of recommendations were developed. The first was generally applicable to all Web pages, and the second addressed the unique concerns of interactive Web-based applications. These recommendations were extensively reviewed and refined by members of the local Web development community. In several cases, compromises may be noted where guidance was tempered to accommodate their diverse concerns. Appendix 10.B presents the general Web page guidelines, and Appendix 10.C provides guidelines for interactive Web applications.

PRACTICAL USE OF TRADITIONAL INTERFACE DESIGN GUIDELINES

In November 1996, we conducted a study in which we compared a sample of recommendations made for various Web products (34 applications and collections of information-only pages) to the 270 traditional recommendations. Of these, only 63 were specifically invoked. Several recommendations were violated fairly frequently (3 were violated by 8 different products). Any 1 application could violate several recommendations, but in fact, violations were distributed among the Web products; the largest number of violations by any 1 application was 22. Table 10.3 shows the recommendations violated most frequently.

A review of Table 10.3 indicates that the guidelines discussed earlier and presented in Appendixes 10.B and 10.C had addressed each of the most frequent observations, although in three cases the guidelines only partially address the observation: information in directly usable format, functionally related data displayed together, and coding used to draw attention. This observation offers some validation for the guidelines development process. Despite the guidelines being communicated through various mass distributions and readily accessible through the intranet, noncompliance has remained a problem. It is impossible to speculate how these data might have differed were guidelines unavailable.

Readers should note the common factor in all the observations listed in Table 10.3: They require knowledge of the user! Because the most basic axiom of

TABLE 10.3
Recommendations Violated Most Frequently From Interactive Web Applications and Web Pages

Recommendation	*# Products in Violation*
Data should be entered in units that are familiar to users.	8
At any given time, pages should present only information essential to the users.	8
Information should be displayed in plain, simple, concise text whenever possible.	8
Functionality should be visible.	7
Information should be presented in a directly usable format that does not require decoding, interpretation, or calculation.	6
When data have a naturally occurring order (e.g., chronological or sequential), it should be reflected in the organization of data fields.	5
White space should be used to create balance and symmetry and lead the eye in the appropriate direction.	4
Horizontal scrolling should be avoided.	4
Functionally related data should be displayed together, on one page.	4
Coding (e.g., color, size, highlighting) should be used to draw attention to critical information, unusual values, changed items, items to be changed, high-priority messages, special areas of the display, and errors of entry.	4

Human Factors is "Know thy user and you are not thy user" (Lund, 1996), it may be concluded that the Human Factors staff must be more involved in the development of Web products. But even in companies where there are no Human Factors practitioners, these data indicate that time should be spent gathering better information about the environment, tasks, and priorities of end users.

We have observed that most violations of Human Factors guidelines have little in common. They seem to draw randomly from the universe of potential interface flaws. This fact poses an interesting problem for the Human Factors Department at Sandia. We have searched for ways to educate developers about principles of interface design to improve usability. The results suggest that we can prescribe guidelines to prevent certain frequent violations of Human Factors principles, but they also suggest that the majority of violations are too frequent and randomly distributed to be addressed by this method. In other words, there is no substitute for incorporating Human Factors into the interface design and testing process.

SUMMARY

We have demonstrated that Web style guide content differs from that found in the traditional Human Factors literature. In general, Web style guides place less emphasis on traditional facets of interface usability, and they often highlight issues such as how authors may overcome the technical problems encountered with distributed, remote access, hypermedia systems. When they do discuss usability, their emphasis reflects the priorities of a hyperlinked environment:

common features across pages (e.g., navigational cues). But why have Web style guides largely neglected Human Factors?

First, it may be argued that although there are overlaps, the WWW introduces unique human-computer interface concerns. This explanation is supported by the emphasis placed on common look and feel in the Web style guides. It is reasonable to suppose that in the hypertext environment of the WWW, where users may readily move between unrelated sites, common look and feel is of greater importance than with stand-alone software applications where differences are less apparent.

As a second explanation for the apparent differences between Web and traditional style guides, Web style guide development has occurred with little rigor and little reference to traditional style guides and principles. The poor agreement between the Web style recommendations and the actual findings of Human Factors reviews of interactive Web-based applications supports this explanation. This was untrue, however, for relatively static Web pages that only presented information or served as waypoints. For these pages, the Web style guides accounted for 64% of findings, which is as good as might be expected (Lowgren & Lauren, 1993; Thovtrup & Nielsen, 1991). This result might be explained by the greater representation of information-only pages than of application pages on the WWW. Inexperience with Web applications may also help to explain the scarcity of application-specific recommendations in Web style guides.

Web style guide authors tended not to reference traditional style guides when creating their own. With the explosive growth of the WWW, it is not surprising that some authors, from diverse backgrounds, would not think to look beyond the WWW for answers to its interface problems. But the appearance of these Web style guides would not be possible without the WWW itself; the ability to publish a personal document and to disseminate it widely is a new phenomenon (December & Randall, 1994).

To be fair, many of the style guides were written for specific audiences (e.g., to assist authors in a particular company). Contrast this with the traditional style guides that were formally sponsored to address specific, but more comprehensive, concerns. In addition, the WWW is still new. Its use to conduct work and business is still being explored; as it matures, it may undergo the same pressures to standardize as did the early graphic computer environments.

Another general conclusion is that there is no substitute in the software development process for Human Factors techniques. Without help from the Human Factors discipline, the number and kind of interface violations that a software developer may unwittingly introduce is huge.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy under Contract DE-AC04-94AL85000. The authors would also like to thank Stephen Crowder of Sandia National Laboratories for his help with statistical analyses.

REFERENCES

December, J., & Randall, N. (1994). *The world wide web unleashed.* Indianapolis, IN: Sams.

Lowgren, J., & Lauren, U. (1993). Supporting the use of guidelines and style guides in professional user interface design. *Interacting With Computers, 5*(4), 385–396.

Lund, A. M. (1996, May). Only nerds call them heuristics: Expert ratings of "rules of thumb" for usability. *Communications Technical Group News*, 1–3.

Thovtrup, H., & Nielsen, J. (1991). Assessing the usability of a user interface standard. In *Conference Proceedings of the Association for Computing Machinery's Special Interest Group on Computer-Human Interaction '91*, (pp. 335–341).

Appendix 10.A: Traditional User Interface Design Recommendations Rated Most Essential by Human Factors Practitioners

1. Data entry should be accomplished through a distinct action.
2. Feedback should be provided about log-on procedures to indicate log-on status.
3. Each color should represent only one category of displayed information.
4. Displayed information should be readable to the degree of accuracy required by the task.
5. Push buttons related to overall window functionality should be located at the bottom of the screen.
6. Axes of line graphs should be clearly labeled with units of measurement, if appropriate.
7. Where users has the option of deleting an entry that is not displayed, they should have the option of viewing the entry before confirming deletion.
8. Information should be presented in a directly usable format that does not require decoding, interpretation, or calculation.
9. Text should be readable from normal viewing distances.
10. When data have a naturally occurring order, it should be reflected in the organization of data fields.
11. Users should be able to distinguish between available and unavailable options.
12. Push button labels should be short and unambiguous.
13. The largest text size should not exceed 10% of vertical display height.
14. Where displayed data exceed a single display, an easy mechanism should be provided to move back and forth through material.
15. Standard scaling conventions should be used: Values on the horizontal axis increase to the right; values on the vertical axis increase going up from the origin; independent variables are plotted against the horizontal axis; dependent variables are plotted against the vertical axis.
16. When users omit data, a message bringing attention to the omission should appear, and either immediate or delayed input of the missing items should be allowed.
17. Users should be required to correct an invalid action and not repeat the entire sequence of actions.
18. Labels should be placed in proximity to their related data fields.
19. Error messages should reflect a user's viewpoint, not a programmer's.
20. User actions should be interpreted in context with previous actions; users should not be required to reenter data.
21. Location of recurring functional groups or items should be the same from panel to panel.

22. Spelling and other common errors should not initiate a transaction different from that intended.
23. A description of the format of specific commands or a list of options should be provided.
24. For fields of five or more characters, the field should be broken into subgroups.
25. Criteria should be established for prioritizing information, and this priority should be utilized in the placement of information.
26. Error messages should indicate the error, the location of the error, and the corrective action.
27. Data fields to be compared character by character should be positioned one above the other, with characters to be compared aligned.

Appendix 10.B: General Web Page Design Guidelines

1. Title: All pages should have a title. The title should appear within title tags and the title should also be displayed on the page either as the highest level heading or incorporated in a graphic or banner. The title displayed on the page and that appearing in title tags should match. The title should be context independent, interpretable by a user who has accessed the page following an unintended route. But, where context-independent titles are not practical (e.g., manuals), other mechanisms may be substituted to provide users with contextual cues. Although it is unnecessary for linked text and associated page titles to match verbatim, there should be sufficient consistency and similarity to assure users they have reached their intended destinations.
2. Signature/contact: All pages should indicate, at a minimum, the name or title (webmaster is acceptable), and E-mail address (phone number if no E-mail) of a contact person. E-mail addresses should be E-mail enabled.
3. Date stamp: All pages should state the date on which the page was last updated. The date data were last updated should be included with query results. The date stamp should appear at the top or bottom of the page, although other placements may be acceptable. With adherence to this guideline, other dates (e.g., issue date for manuals) may be used at a page owner's discretion.
4. Current links: All links should be active, although nonhighlighted placeholders are acceptable where sites are under construction and will be available in a short time. Links to empty pages with "Under Construction" or similar banners should be avoided. Where pages contain links to other sites (both in and out of the Sandia National Laboratories intranet), the links should be periodically checked; when inactive links are found, updates should be made in a timely manner.
5. Show file size: Warnings should be offered if a static, non-HTML file (as opposed to dynamic files such as query results) requires an

inordinate time to download. Design solutions that do not have the overall effect of unacceptably cluttering pages should be sought. It is generally recommended that the file size be given for all non-HTML files 65 k or larger. This guideline does not necessarily encompass links to HTML files, although in developing HTML files, developers are encouraged to be aware of the effect of file size on end users.

6. Alert to unexpected properties of links: If selecting a link executes actions other than calling another Web page (e.g., launches an application, Telnet, FTP, downloads a large file), users should be alerted. For example, "Account/Password Form (displays the form as an Adobe Acrobat.pdf file)", or "Policy Statement (a 75 K Microsoft Word document)".
7. Text Size: Unusually small text (including text contained in graphics and tables) should be avoided. As general guidance, the U.S. Department of Defense Human-Computer Interface Style Guide (30 Sep, 1992) calls for a minimum text size of approximately 3 mm, whereas the International Standards Organization calls for greater than 16 arc min. Although text appears different across platforms and with different fonts, this guidance generally translates to Hypertext Markup Language (HTML) no smaller than H6 with H5 preferred, and for word processors this translates to fonts no smaller than 10 points with 12 points preferred.
8. Labeling: All labels, whether linked text, buttons, or graphic features, should be clear and concise. Labels used for graphics, tables, or data entry fields should appear in proximity to the object described by the label. Linked text should be descriptive and meaningful. The use of "Click here" and similar labeling of links should generally be avoided, although where necessary, instructions for the use of image maps and other less straightforward interactive features are acceptable. Furthermore, labels should use terminology familiar to the intended user or reader.
9. Platform consistency: The layout of pages should be essentially the same (including no chopped-off features) when displayed on personal computer, Macintosh or Unix platforms, with browser in the default setting (out-of-the-box configuration). Platform consistency may be unnecessary for pages intended for a small, well-known audience (e.g., project personnel, all using Silicon Graphics).
10. Clutter: The layout of pages should be balanced. Various mechanisms (paragraph breaks, blank lines, horizontal bars) should be used to minimize clutter, although such mechanisms should not be used to the extent that they themselves serve as a source of excessive clutter.
11. Page size: Generally, HTML file length (or Web page) should be limited to a size that allows the entire page to be viewed in three, preferably two, consecutive screens (14-inch monitor), with browser in default settings. The intent is to avoid pages that require excessive

scrolling; other solutions to this general goal are acceptable. In some cases (e.g., technical documents that need to be accessed in their entirety, other long documents that are meant to be printed, or pages that are infrequently used), longer files are acceptable.

12. Information layout: Generally, the order in which information is presented should reflect the importance, priority, and frequency with which the information is accessed; strict ordering may be abandoned if necessary to attain the most meaningful organization of information.
13. Backgrounds: The use of complex background patterns or background colors that degrade the readability of page contents should be avoided.
14. Functionality visible: All links, buttons, or other functional elements of pages should be readily apparent to users.

Appendix 10.C: Interactive Web Applications Guidelines

1. Data acceptance or rejection: The application should provide users positive feedback (the absence of a system response is not sufficient feedback) that data submitted to the system have been either accepted or rejected.
2. Log-on prompt: Applications that require log-on procedures should provide prompts to explain any log-on failures.
3. Help: On-line Help should be available and accessible from the application. Furthermore, Help should be organized to facilitate users obtaining answers to questions about use of the application; users should be able to select from topics and, if appropriate, levels of Help.
4. Omissions: Applications allowing users to submit data should inform them of required data entries they may have omitted in preparing their data submission.
5. Error correction: Applications should provide an easy mechanism for users to correct erroneous data entries and should alert users to the location of errors. Furthermore, users should be required to re-enter only erroneous entries.
6. Meaningful messages: Messages should be meaningful and stated in terms familiar to users; they should tell users what steps are necessary to proceed with their transactions.
7. Positive indication function actuation: Applications should positively indicate (absence of a system response is not sufficient) when application functions have been actuated.
8. User confirmation: The application should require users to confirm any action that could have a potentially destructive effect (e.g., erasing or overwriting data).
9. Minimize user action: Data entry mechanisms should be designed to logically minimize actions required of users, including not requiring users to re-enter data, automatically entering known data, minimizing keystrokes, avoiding short-term memory loads.

10. Data entry: Complex rules for data entry should be avoided. Data entry fields should be as flexible as may be reasonably achieved, and should accept data in multiple formats.
11. Required or optional fields: Cues should be provided to aid users in distinguishing required from optional data entry fields.
12. Process abort: Where practical, applications should inform users when, for various reasons, processes have been aborted.
13. Data entry fields: Data entry fields should be sufficiently large to accept any reasonable user entry.
14. Sequencing: Items should generally be ordered in accordance with logical or established task sequences.

11 Page Design Guidelines Developed Through Usability Testing

José A. Borges
Israel Morales
Néstor J. Rodríguez
University of Puerto Rico at Mayagüez

The World Wide Web (WWW) has become a popular means for publishing information. A large number of information repositories (Web sites) already exist, and new ones are rapidly being created. Most pages on Web repositories provide elements that allow users (readers) to interact with them. Thus, people designing pages for the WWW are actually designing user interfaces (UIs).

The proliferation of pages with poor usability suggests that most designers of Web pages have little knowledge of UI design and usability engineering. This serious problem needs to be addressed, as pages with poor usability can have these negative effects:

- Users are frustrated, because of an inability to find the information sought, disorganized pages and confusing information, pages under construction and disconnected links, the lack of navigation support, and other problems.
- Exploration is discouraged, because of barriers imposed by the poorly designed interface and users' lack of trust or faith in the site.
- Time is wasted, because of disorganized pages, misleading link names, long pages, and long download times.
- Internet traffic is increased, because the problems just mentioned not only affect the use of a particular site, but are also responsible for much unnecessary traffic on the Internet.

Web pages with poor usability can also have negative effects for site owners: They can result in a reduction of visits to the site, negative feedback from frustrated users, and a negative image for the site. Thus, from all points of view, it is important to develop practical methodologies for designing usable Web pages.

The topic of Web design has become very popular in recent years, manifested by the increasing number of books, courses, tutorials, and articles on the subject. Many guidelines for designing Web pages can be found on the WWW. Among the most comprehensive are Guide to Web Style, by Rick Levine from Sun Microsystems (Levine, 1995); the Web Style Manual, by Patrick Lynch from the Yale Center for Advanced Instructional Media (Lynch, 1995); and

Elements of Web Design, by Ben Benjamin from The Computer Network (Benjamin, 1996). These works are complete and extensive; most are based on personal experiences and observations and are not supported by formal experimentation. Shorter articles on this topic include the papers by Hagan Heller and David Rivers (Heller & Rivers, 1996) and Tim Comber (Comber, 1995).

Various methods for designing pages with good usability are suggested in two studies conducted by Jakob Nielsen (Nielsen, 1995; Nielsen & Sano, 1994), but some of these methods are not practical for most page designers because they require user testing and are time consuming. Also, most designers of Web pages lack the appropriate facilities, resources, and usability engineering expertise to effectively apply some of the recommendations.

A more practical alternative for most page designers is a set of simple guidelines for designing Web pages, guidelines that avoid the common mistakes found in many pages today and that promote usability. These guidelines should be based on fundamental principles of UI design, but should be simple enough for nonexperts to understand and apply. The goal of our work has been to develop a set of guidelines that meet the following criteria: The guidelines should be short, simple, and practical; in other words, they should be easy to read and understand. Users should not need to know UI design and usability engineering to apply the guidelines, and they should be supported by usability evaluation and testing. The following sections describe the development process and usability evaluations performed to generate a set of guidelines for designing usable Web pages.

HEURISTIC EVALUATION OF WEB SITES

The design of a good UI usually requires various usability evaluation methods. Among these, the heuristics evaluation method (Nielsen, 1994) is the least expensive, yet one of the most useful. In our quest for a good set of guidelines for designing Web pages, we started by conducting a heuristic evaluation of pages of existing Web sites. Because of the diversity of existing repositories, we thought it necessary to begin the study with a group of Web sites with similar characteristics. Thus, we decided to start the study with a heuristic evaluation of university and college Web sites.

A sample of 10 repositories was randomly selected from a listing of more than 1000 colleges and universities (DeMello, 1997). Home pages (the first page of a site) and secondary pages (the other pages) of each site were evaluated by three evaluators. The evaluation was based on a set of 10 usability heuristics (guidelines) for designing UIs (Nielsen, 1994).

The heuristic evaluation detected a variety of usability problems, most of which were related to the following three heuristics developed by Jakob Nielsen.

> *Aesthetic and minimalist design*: Dialogues should not contain irrelevant or rarely needed information. Every extra unit of information in a dialogue

competes with the relevant units of information and diminishes their relative visibility.

Match between system and the real world: The system should speak users' language, with words, phrases, and concepts familiar to users, rather than use system-oriented terms. By following real-world conventions, information appears in a natural and logical order.

Consistency and standards: Users should not have to wonder whether different words, situations, or actions mean the same thing. Platform conventions should be followed.

Dialogue Problems

Most problems encountered were related to the *aesthetic and minimalist design* heuristic. Some pages evaluated had headers that took a substantial portion of a page's visible area. By reducing the visible area available for placing information, designers forced users to scroll more than usual. In some cases, headers and footers were not clearly separated from the body of a page; these pages were perceived as disorganized because the information was not properly grouped. Because headers are mostly used to emphasize a page's theme, they must be clearly separated from the body of the page to get users' attention. Footers usually provide information that is not directly related to the content of a page. As they serve a different purpose, they must be clearly separated from the body of the page.

Various pages had links with explanatory comments attached. The intent of these comments was to provide a notion of the content of the pages that the links pointed to. In most cases, this practice is unnecessary because it adds redundant information to a page. The name of a link should be sufficient to provide a hint of the page it points to. Another undesirable practice was that of making a link every time the same keyword was mentioned in a segment of text. We refer to this practice as linking mania. This practice adds redundant links that compete with other links and with the page's relevant information.

Some dialogue problems encountered were unique to the home pages of the repositories (the primary pages). The home pages are the gateway of a Web site; their main function is to provide an idea of the most important aspects of the site and to direct users to them. Because these pages are the most frequently accessed pages of a Web site, they should be designed in such a way that users can get a good idea of the main aspects of the site and quickly move to the path to the desired information.

Some home pages included one or two paragraphs about the organization. Although intended to advertise the organization, this information could be irrelevant for most users, especially those who visit the page regularly. In addition, the inclusion of such information increases the size of the page and forces users to scroll. As a result, users need a long time to get a feel for the Web site and to navigate. According to usability studies conducted by Jakob Nielsen (Nielsen, 1996), users do not want to scroll and do not want to read.

Some home pages were not well organized and thus, did not provide a good idea of a site's most important aspects. One problem encountered was the disorganized mixing of links to primary and secondary topics. Another problem was the inclusion of links to nonprimary pages of the site.

Language Problems

We commonly found links with names that did not provide a good hint about the page they pointed to. Some link names included words that were not commonly understood by most users. Such link names can introduce navigation problems because users could follow the wrong links. This action wastes user time and causes unnecessary traffic on the Internet. A similar problem results when images or icons are used to link to other pages. If these images do not provide a distinctive feature of the page they point to, users have no idea of the information they are accessing.

Consistency Problems

Most consistency problems were related to footers and navigation elements. Some sites did not make consistent use of footers and navigation elements across the site.

Other Problems

Two additional problems are worth mentioning. First, some links pointed to pages that did not exist. Following one of these links would obviously waste user's time. The second problem was related to the reproduction of hard copies of the pages. Some pages had images with colors that did not reproduce well on black and white printouts, and valuable information was lost when the pages were printed.

Useful Features

During the heuristic evaluation, the evaluators noticed two useful features on some pages. One was the inclusion of a link to an index of pages for the site. This feature can be very useful because it allows quick access to the pages of the site without much navigation. The other feature was the inclusion of a footer with information and links related to the pages. Such footers usually indicated information such as the last date the page was modified, an E-mail address for sending comments about the page and the Internet address of the page.

On the basis of the usability problems encountered from the heuristic evaluation of college and university Web sites, we compiled a list of guidelines for designing Web pages (see Table 11.1).

USER TESTING OF GUIDELINES FOR DESIGNING WEB PAGES

To evaluate the proposed guidelines, we conducted an experiment with two versions of the home pages of three Web sites randomly selected from the 10

TABLE 11.1
Proposed Guidelines for Designing Usable Web Pages

For Any Page

1. Headers should not take more than 25% of a letter-size page.
2. Headers and footers should be clearly separated from the body of the page. (One way of achieving this is by placing bold lines or horizontal rules between them and the body.)
3. Names of links should be concise and provide a hint of the content of the page they link to. (Avoid using technical words such as servers, links, web server)
4. Avoid adding explanatory comments to textual links.
5. Avoid linking mania (making a link every time a keyword of a page is mentioned in a text).
6. Verify that links connect to existing pages.
7. Linking icons should have a distinctive feature of the page they link to.
8. Maintain consistency when using icons. The same icon should be used for the same purpose.
9. Colors should be selected so that pages can be clearly displayed and reproduced on black and white displays and printers.
10. It is desirable to include the date the page was last modified, the mail address of the person who maintains the page, and the Uniform Resource Locator (URL) address of the page on a footer.

For the Home Page of the Repository

11. Descriptive information about the institution should be placed just below the header and kept to a minimum (a link to a secondary page is preferable).
12. Pages should not be overcrowded with links.
13. Pages should be short (about a letter-size page).
14. Links should be to primary aspects or characteristics of the institution. Textual information should be left for secondary pages.
15. Organize links as primary and secondary topics.
16. Links to resources or other repositories on the Internet should be placed on a secondary page. (This page should be reached with a link on the primary page.)
17. A more extensive index of links, properly grouped, can be provided on a secondary page for fast access to a wide range of the institution's repositories.

Note. From "Guidelines for Designing Usable World Wide Web Pages," by J. A. Borges, I. Morales, & N. J. Rodríguez, 1996, *Conference Companion of the Computer-Human Interaction Conference,* pp. 277–278. Copyright © 1996 by ACM.

college and university Web sites evaluated (Borges, Morales, & Rodríguez 1996). The two versions were the original home page and a redesigned version of the original home page. The redesigned versions were developed by three people with Web page design experience. Each designer was assigned a different home page and was given a copy of the proposed guidelines, with instructions to redesign the home pages, without any constraint except to apply the guidelines provided. No explanation was given about the guidelines.

Ten users with different levels of Web experience were asked to perform five tasks using the original versions, and another ten users used the redesigned versions. Tasks 1, 2, and 3 required users to select a link at the top, middle, and bottom of the home page, respectively. The fourth task requested users find the home page of the Computer Science Department. The fifth task requested users to find information about the university's library resources.

A summary of the results is presented in Table 11.2. The times indicated in the table correspond to the average time that the 10 users took to complete each task. The results of the test indicated that the average time to complete most tasks using the redesigned home pages was significantly reduced. In three cases users needed more time to complete the task with the redesigned page. To get a better understanding of these results, we analyzed both versions of the three home pages according to their compliance to the guidelines.

The original version of Home Page A was too long (about two letter-size pages long), it was overcrowded with links and explanatory comments, and it used approximately one fourth of its space to provide information not related to the site. The redesigned version was a short page (about half a letter-size page) that improved all these aspects and thus was more in compliance with the guidelines. These improvements in the redesigned version helped improve the usability of the home page.

The original version of Home Page B was too long (about two letter-size pages). It had a long header (half a letter-size page) and many links with explanatory comments. The redesigned version was a shorter page (about half a letter-size page) with a smaller header and fewer links (none with explanatory comments). The redesigned version did not include a link to an index page (a page consisting of links to most of the pages of the site) that was present in the original version. This link was used by several users of the original home page to speed up the information search on Tasks 4 and 5. We believe that the omission of this

TABLE 11.2
Results of Usability Test of Web Page Design Guidelines

	Home Page A			*Home Page B*			*Home Page C*		
	Average Time (seconds)		*Percent Reduction*	*Average Time (seconds)*		*Percent Reduction*	*Average Time (seconds)*		*Percent Reduction*
	Original	*Re-designed*		*Original*	*Re-designed*		*Original*	*Re-designed*	
Task 1	24.0	5.4	+ 78	11.0	7.2	+ 35	7.0	4.3	+ 39
Task 2	13.1	2.3	+ 82	12.2	3.7	+ 70	4.0	18.5	- 363
Task 3	11.9	3.8	+ 68	20.4	8.2	+ 60	4.4	7.1	- 61
Task 4	31.9	19.2	+ 40	17.9	29.2	- 63	21.9	5.4	+ 75
Task 5	35.8	30.1	+ 16	13.4	6.8	+ 49	11.2	6.9	+ 38

Note. From "Guidelines for Designing Usable World Wide Web Pages," by J. A. Borges, I. Morales, & N. J. Rodríguez, 1996, *Conference Companion of the Computer-Human Interaction Conference,* pp. 277–278. Copyright © 1996 by ACM.

link in the redesigned version was responsible for the increase in the average time to complete Task 4.

The original version of Home Page C was a short page (about half a letter-size page) with very few links. This page complied with most of the guidelines. The redesigned version was a long page (longer than a letter-size page) with a longer header (about one third of a letter-size page) and more links. The increase in the average time to complete Tasks 2 and 3 for the redesigned version was due mostly to the increased page and header sizes. But an additional factor contributed to the large time difference that resulted for Task 2. The redesigned home page featured four large arrow-like icons below the header. Only two of these icons were visible on the browser window at the beginning of the test. It seems that users perceived these icons as navigation buttons that marked the end of the page because most of them did not attempt to scroll the page, a move was necessary to reach the link that allowed them to complete the task. The page and header sizes were not major factors for Tasks 4 and 5 because the redesigned page featured, in a prominent place, a link to a local search facility that many users used to complete the task. This link was also available in the original home page, but was not easily visible.

Our evaluation of the redesigned home pages showed that none of them properly complied with guideline 14 of the proposed guidelines. In an interview with the designers of the redesigned pages, they stated that their selection of topics and keywords was greatly influenced by the original home page assigned to them. Another significant observation was that guideline 17 was discarded by the three designers. From the interview with the designers, we concluded that this guideline was not clearly understood.

HEURISTIC EVALUATION OF COMMERCIAL WEB SITES

The experiment described in the previous section demonstrated that designers of Web pages can improve the usability of home pages by applying the proposed guidelines, but a wider sample and variety of sites needed to be evaluated to refine and produce a more complete set of guidelines. With this objective in mind, we decided to conduct a similar study with commercial sites. Twenty sites were randomly selected for the study. The selected sites were of a heterogeneous business nature (book publishing companies, a carmaker, a pet shop, a communications company, entertainment companies, specialty stores, food stores, a money-transferring company, an electronic magazine, and computer shopping companies). None of the sites had previously been visited by any of the three evaluators who participated in the study. Again, the heuristic evaluation was based on the usability heuristics developed by Nielsen (Nielsen, 1994).

In general, the pages of the commercial sites made more use of artistic images and were more attractive than the pages of the college sites. In terms of usability problems, however, they were similar to the college and university home

pages. A significant number of home pages were not well organized and were crowded with text. Almost half of the home pages were relatively long.

Some minor differences from the college pages are worth noting. For example, on the commercial sites, the headers of some of the home pages did not clearly identify the organization. There were more cases of image maps that did not provide an indication that they contained links, and there was no evidence of linking mania or of pages crowded with links Because these last two problems were not very common on the college sites, guidelines 5 and 12 did not address a significant problem.

As the commercial Web sites exhibited usability problems similar to the college and university home pages, their pages could probably be improved by applying the proposed guidelines to their design.

IMPROVING NAVIGATION

One of the most critical aspects of browsing pages on the WWW is navigation. Because of the way most Web sites are structured, users spend significant time moving from page to page in search of information. Inappropriate and misleading link names, confusing and disorganized pages, and other usability problems contribute to the disorientation of users who can become lost in hyperspace.

From our heuristic evaluations of Web pages, we identified three page design practices that could potentially improve navigation on a Web site. These practices are the selection of appropriate link names, the adoption of an index of links to pages or topics, and the use of navigation buttons. In the following sections, we discuss how these practices can improve navigation and how to integrate them into the page design process.

Selecting Appropriate Names for Links

The selection of names for links is a critical aspect in the design of Web pages. In most cases, the name of a link is a user's only hint about the content of the page the link points to. If the name is misinterpreted, a user could reach a useless page in the search for information. This situation results in a waste of time and causes unnecessary traffic on the Internet. Thus, it is important that this problem be minimized in the design of Web pages.

Page designers can minimize the link-naming problem following guidelines for selecting appropriate link names. To produce such guidelines, we found it necessary to get a better understanding of the nature of the problem. We conducted a study on how well link names reflect the content of a page (Rodríguez, Borges, & Morales, 1997). The study consisted of asking users to indicate what they thought was the content of the page pointed to by a particular link. Fifty links, randomly selected from 10 commercial home pages (also randomly selected), were used for the study. Sixteen people participated in the study. All were given a hard copy of the 10 home pages with the selected links

identified with a number and they were asked to write, on a response sheet, what they thought was the content of the page pointed to by each of the 50 links.

Each response was classified according to the clue that the link name provided about the content of the page it pointed to. The following four categories were used to classify the responses:

None: The link name did not provide an idea of the content of the page.
Wrong: The link name provided the wrong idea about the content of the page.
Fair: The link name provided some idea about the content of the page.
Good: The link name provided a good idea about the content of the page.

The category assigned to each response was determined by a consensus of the subjective opinions of three experimenters. The experimenters' opinions were based on their analysis of the content of the 50 pages pointed to by the links selected for the study.

In Table 11.3, we present a breakdown of all the responses by category. The responses indicate that in most cases the names assigned to a link by page designers provided a good idea or some idea about the content of the pages the links pointed to. A significant number of links, however, did not provide a good hint. In approximately one fourth of the cases, the link names suggested a wrong idea about the content of a page. If these links were followed, user's could waste time. We also found that in approximately 9% of the cases, the link names did not provide a clue about the content of the page. These link names could discourage exploration or could result in a waste of time, if followed.

To get a better understanding of why some link names provide poor clues, we analyzed the pages pointed to by the 10 link names that provided the worst clues. From this analysis, we found that in most cases the name selected by the page designer did not reflect the content of the page. In other cases, the names were too general and did not provide enough information for a person to at least guess at the page content. Some names were not easily understood, and most people would have had no idea about the content of the page pointed to by the link.

TABLE 11.3
Clues Provided by Link Names About the Content of the Pages They Pointed to

Category	*Percentage of Cases*
None	9.3
Wrong	25.3
Fair	30.1
Good	35.3

On the basis of the previous analysis, we found it necessary to provide a more specific guideline for naming links. We propose the following:

- Link names should reflect the content of the page they point to.
- Link names should be specific. For example, it is more meaningful to use Crosswords and Comics instead of Recreation and Deadlines instead of Important Stuff.
- Avoid names that are not commonly understood such as technical words and terms with in-house or local meanings. For example, it is more meaningful to use Industrial Affiliates Program instead of IAP and Computer and Communications Technologies instead of Telematics.

Using Indices

An index of pages or topics of a Web site can be very useful: It allows quick access to the pages of the site without much navigation. During user testing of the original guidelines for designing Web pages, we noticed that once users discovered an index page, they would frequently return to the page while searching for information on the site.

To study the effectiveness of indexes, we conducted two experiments involving index pages. In the first experiment, we compared the effectiveness of two types of index pages. One was a linear list of links to different pages of a site (see Fig. 11.1). The header of the page featured the letters of the alphabet as links to each of the groups. The header was static (did not scroll). The other index page consisted of the same list of links used on the alphabetical index page, but organized like a table of contents (see Fig. 11.2). The header provided links to the main topics of the site (the topics that appear in the home page of the site).

Two groups of 10 Web users participated in the study. Both groups were asked to perform 15 tasks that required them to look for specific information from different pages of a Web site. These pages were randomly selected from the list of links of the index page. One group performed the tasks using the alphabetical index, and the other group used the table of content index. All the tasks were started on the particular index page. No instructions were provided to users on how to use the index page. The time users took to select the index page link that allowed them to complete the task was recorded. Only the time spent by the users on the index page was counted.

The results of the experiment are presented on Table 11.4. The times presented in this table correspond to the average time it took users to find a link on each index page for each task. In 13 out of 15 tasks, the average completion time using the alphabetical index was lower than that achieved by the users of the table of content index. In 9 tasks, users of the table of content index took, on average, more than twice the time that users of the alphabetical index needed to complete the tasks. Thus, it is evident that in terms of access time, the alphabetical index is superior to the table of content index.

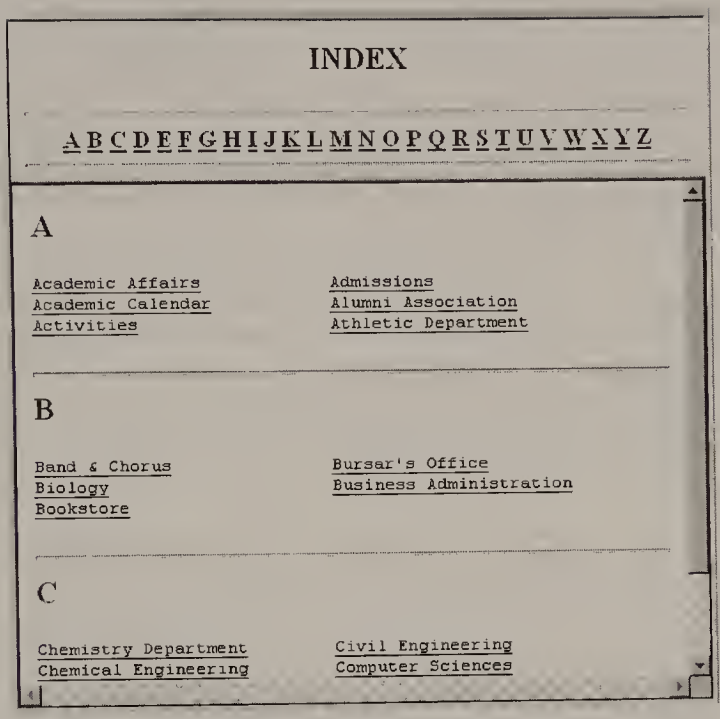


FIG. 11.1. Alphabetical index page.

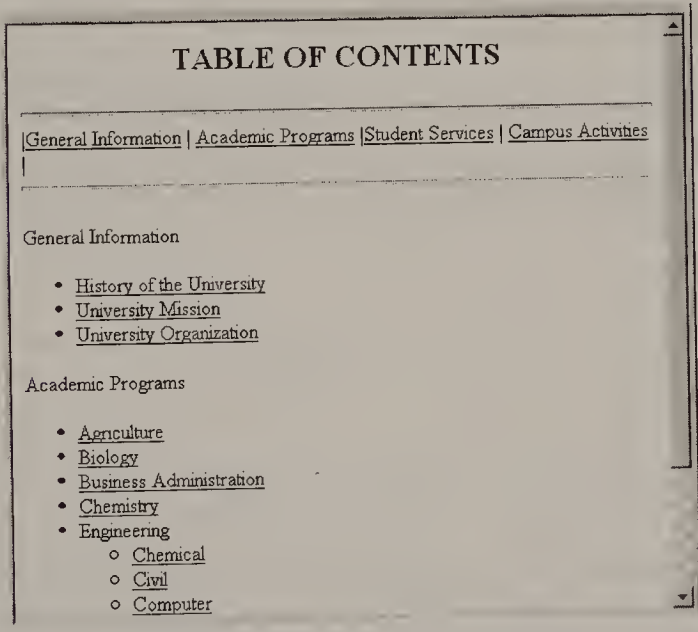


FIG. 11.2. Table of content index page.

TABLE 11.4
Comparison of Average Times to Find Links on an Alphabetical Index Versus a Table of Contents Index

	Alphabetical Index *(average time in seconds)*	*Table of Content Index* *(average time in seconds)*	*Time Ratio* *(Table of Content /Alphabetical)*
Task 1	22	75	3.41
Task 2	16	46	2.86
Task 3	28	79	2.82
Task 4	14	48	3.43
Task 5	18	10	.56
Task 6	42	94	2.24
Task 7	19	74	3.89
Task 8	20	27	1.35
Task 9	23	54	2.35
Task 10	25	26	1.04
Task 11	21	12	.57
Task 12	17	40	2.35
Task 13	10	18	1.80
Task 14	12	42	3.50
Task 15	17	30	1.76

A second experiment was conducted to determine whether the adoption of an index page can improve the navigation of a Web site (Rodríguez, Borges, & Morales, 1997). Two groups of Web users participated in the experiment. These users were different from the ones that participated in the previous index experiment. One group was asked to perform 15 tasks on a Web site. The other group was asked to perform the same 15 tasks on a modified version of the same Web site. The tasks were the same tasks used in the alphabetical and table of contents index test. The modified version of the site had an alphabetical index of pages of the site. A link to this index page was provided on the home page of the site and on the pages pointed to by the links of the home page. To prevent any bias toward the use of the index page, users of the modified version were not told about the existence of the index page. All the tasks began on the home page of the site. The time to complete the tasks (excluding page download time) was measured.

The results of the experiment are presented on Table 11.5. The times presented in this table correspond to the average times it took users to complete the task on each site version. The results indicate that, on average, users of the site with an index were able to complete most tasks in a shorter time than were users of the site without an index. In 11 out of 15 tasks, the average completion time for users of the site with an index was lower than the average completion time users

of the site without an index. The tasks that users of the site without an index completed in a shorter average time were 4 of the first 5 tasks of the experiment. During these first 5 tasks, many users of the site with an index were not aware of the existence of the index page and thus could not benefit from it. After the fifth task, most users were aware of the index page and used it frequently to speed navigation of the site. Thus, a fair comparison of the two versions of the site should be based on the results of the last 10 tasks. It is evident from the results of the last 10 tasks that the adoption of an alphabetical index can improve navigation on the WWW because it helps users find information faster.

From the experiment, we found additional evidence in support of the incorporation of an index page. There were 23 cases in which a user of the site without an index page could not complete a task. This contrasts with only 2 cases on the site with an index page. These results suggest that the incorporation of an index page results in a higher degree of success in the search for information on a Web site.

Using Navigation Buttons

Navigation buttons have the potential for improving navigation on a Web site because they facilitate a quick move to key pages of the site. From the indexes studies, we observed that when users browse a site, they tend to revisit the home page of the site. In most cases, they do this by consecutive activation of the Back button of the browser. This navigation method is ineffective if the Back button must be activated more than once because users revisit pages that serve no purpose in the search for information. A faster way to the home page can be achieved by providing a button or link on each page that links to the home page. By selecting this button or link, users immediately get back to the home page without having to revisit unnecessary pages.

The experiment demonstrated that when an index page is available, users access it frequently; thus designers must include an index button on each page of a site. Other buttons, such as a predecessor button (that links to the page that is the logical parent of the current page) and a top button (that links to the top of the current page), could be useful, but we have not conducted studies that support their usability.

SUMMARY

The study of commercial Web sites and the navigation studies previously mentioned, have provided information about the design of usable Web pages. This new knowledge called for a revision of the original guidelines presented in (Borges, Morales, & Rodríguez, 1996). The revised set of guidelines is presented in Table 11.6. These new guidelines are expressed in simple terms so that any page designer can apply them. They are helpful to most page designers because they allow them to design usable Web pages without spending too much time reading literature about page design. The guidelines fit on one page and take only

a few minutes to understand. The guidelines are also useful for conducting heuristic evaluations of existing pages to detect usability problems.

As technology advances, new elements are integrated in the design of Web pages. This process requires a constant review of page design guidelines. For this reason, we have developed a Web site to maintain the latest version of the guidelines and to illustrate the use of the guidelines with examples (Rodríguez, 1996). In the meantime, the proposed guidelines could help the new breed of interface designers (Web page designers) apply the principles of human-computer interaction and usability engineering in their designs (the Web pages).

ACKNOWLEDGEMENTS

We thank Ana R. Marrero, Julio Ortega, and Angel L. Ramos for developing the modified pages used in the studies of indexes and for assisting in the realization of some of the usability studies. We also thank Manuel A. Pérez and Domingo Rodríguez for reviewing this article.

TABLE 11.5
Comparison of Average Times to Complete Search Tasks on Web Sites With and Without Index Pages

	Site A (With Index Page) *(average time in seconds)*	*Site B (Without Index Page)* *(average time in seconds)*	*Time Ratio* *(Site B/Site A)*
Task 1	52	12	.23
Task 2	27	16	.59
Task 3	46	41	.89
Task 4	13	25	1.90
Task 5	15	12	.80
Task 6	63	107	1.70
Task 7	64	99	1.55
Task 8	19	91	4.79
Task 9	24	31	1.29
Task 10	17	24	1.41
Task 11	18	27	1.50
Task 12	27	118	4.37
Task 13	16	115	7.19
Task 14	18	179	9.94
Task 15	23	47	2.04

Note. From "Page design guidelines for improving WWW navigation," by N. J. Rodríguez, J. A. Borges, & 1. Morales, 1997, *Proceedings of the Seventh International Conference on Human-Computer Interaction.* Copyright © 1997.

TABLE 11.6
Revised Set of Guidelines for Designing Usable Web Pages

For Any Page

1. Headers should not take more than 25% of a letter-size page.
   Headers and footers should be clearly separated from the body of the page.
2. Names of links:
3. - Link names should reflect the content of the page they point to.
   - Link names should be specific. For example, it is more meaningful to use Crosswords and Comics instead of Recreation and Deadlines instead of Important Stuff.
   - Avoid names that are not commonly understood such as technical words and terms with in-house or local meanings. For example, it is more meaningful to use Industrial Affiliates Program instead of IAP and Computer and Communications Technologies instead of Telematics.

   Linking icons should have a distinctive feature or phrase of the page they link to.
4. Adding explanatory comments to textual links should be avoided.
   Links that connect to existing pages should be verified.
5. Icon use should be consistent. The same icon should be used for the same purpose.
6. Colors should be selected so that pages can be clearly displayed and reproduced on
7. black and white displays and printers.
   An index page should be provided on a Web site. This page should be constructed as
8. a alphabetical list of links to pages or topics of the site. The links or topics should be grouped and listed in alphabetical order (one group for each letter of the alphabet). A
9. link to each group should be provided in the header of the index page (A B C ... Z). The header should remain visible on the screen while the list of links is scrolled.
   Navigation buttons should be provided on all the pages of the site. These buttons should be placed below the header of the page and above the footers of pages that require scrolling for viewing. At least the following buttons should be provided: one that links to the index page and one that links to the home page of the site. The
10. buttons should be placed close to the right side of the page.
    It is desirable to include the date that the page was last modified, the mail address of
11. the person who maintains the page, and the Internet address of the page on a footer.

For the Home Page of the Repository

12. Textual information should be minimized.
13. Any information not directly related to the organization should be placed outside the body of the page (the body of the page is the area between the
14. header and the footer).
15. Pages should be short (should not require scrolling).
    Links should be to primary aspects or characteristics of the organization.

REFERENCES

Benjamin, B. (1996). Elements of Web design, CNET (The Computer Network) Inc., <http://www.cnet.com/Content/Features/Howto/Design/> (accessed May 21, 1997).

Borges, J. A., Morales, I., & Rodríguez, N. J. (1996, April). Guidelines for designing usable world wide web pages. *Conference companion of the computer-human interaction conference*, (pp. 277–278). Vancouver, BC, Canada: Association for Computing and Machinery.

Comber, T. (1995, April). Building usable web pages: An HCI perspective. In *Proceedings of the First Australian WWW Conference*. Lismore, New South Wales, Australia: Norsearch. <http://elmo.scu.edu.au/sponsored/ausweb/ausweb95/papers/hypertext/comber/> (accessed May 21, 1997).

DeMello, C. (1997). College and university home pages. <http://www.mit.edu:8001/people/cdemello/univ.html> (accessed May 16, 1997).

Heller, H., & Rivers, D. (1996, March). So you wanna design for the Web. *Interactions*, *3*(2), 19–23.

Levine, R. (1995). Sun on the Net – Guide to Web style. <http://www.sun.com/style guide/> (accessed May 21, 1997).

Lynch, P. J. (1995). Web style manual. Yale Center for Advanced Instructional Media, <http://info.med.yale.edu/caim/stylemanual/Manual-1.html> (accessed May 21, 1997).

Nielsen, J. (1994). Heuristic evaluation. In J. Nielsen & R. L. Mack (Eds.), *Usability inspection methods* (pp. 25–62). New York: Wiley.

Nielsen, J. (1995). A home-page overhaul using other web sites. *IEEE Software*, *12*(3), 75–78.

Nielsen, J. (1996). Interface design for Sun's WWW site. <http://www.sun.com/sun-on-net/uidesign/> (accessed May 21, 1997).

Nielsen, J., & Sano, D. (1994, October). SunWeb: User interface design for Sun Microsystem's internal web. In *Proceedings of the Second World Wide Web Conference: Mosaic and the Web*. Chicago, IL. <http://www.ncsa.uiuc.edu/SDG/IT94/Proceedings/HCI/nielsen/sunweb.html> (accessed May 21, 1997).

Rodríguez, N. J. (1997). Guidelines for designing usable WWW pages. <http://exodo.upr.clu.edu/rumhp/hci/wwwdesign.html> (accessed May 21, 1997).

Rodríguez, N. J., Borges J. A., & Morales, I. (1997, August). Page design guidelines for improving WWW navigation. In *Proceedings of the Seventh International Conference on Human-Computer Interaction*. San Fransico, CA: Elsevier.

12 Human Factors Methodology for Designing Web Sites

Pawan Vora
U S WEST Communications

In the last few years, there has been a tremendous growth on the World Wide Web (WWW)—both on the Internet and on corporate intranets. The Netcraft Web Server Survey identified over 1,000,000 Web sites on the Internet in April 1997, as compared to about 19,000 in July 1995 (see http://www.netcraft.com). Similar growth patterns have been observed on corporate intranets. According to the Forrester Research July 1996 report, two-thirds of the Fortune 1000 companies have an intranet. And according to International Data Corporation (IDC), intranet servers can outsell Internet servers 4.6 million to 440,000 at the turn of the century.

There are several reasons for such an explosive growth on the WWW: ease of publishing; relatively low cost of having a Web presence; low training needs because of the simple, link-based navigational model; ease in updating and distributing information; platform independence; and the potential of reaching a wide audience.

Such seductive benefits of the WWW have led to an unprecedented rush to have a Web presence, but, many Web sites have failed to realize their potential because they disregard users' needs and requirements (Sano, 1996). Evidence of this disregard is obvious in the pervasiveness of the following usability-compromising situations:

Slow response time or speed: Georgia Institute of Technology's Graphic, Visualization, and Usability Center's (GVU) 6th WWW User Survey found that speed continued to be the number one problem of Web users (GVU, 1996).

Difficulty finding information: Although the link-based navigational model of the WWW is easy to understand, accessing information on it has become increasingly difficult. GVU's 6th WWW User Survey reported that after speed, the next big problem was "finding known information" (34.09%) and being able to find pages already visited (13.41%).

Opaqueness of technology: The technology is no longer transparent to users, who now need to understand terms such as plug-ins, helper applications, Shockwave, PDF, Java, JavaScript, ActiveX, and so on, to take full advantage of the WWW.

Poor design and implementation: These are evident in frequent encounters with unimportant, outdated, and useless information; dangling or broken links; At

Work or Under Construction signs; inconsistencies in the design of Web pages; and convoluted navigation to get to desired information.

There are several reasons for the current state of affairs. The main reason appears to be the perceived need to have a Web presence: Many organizations rush to be on the WWW without spending sufficient time to understand the needs of potential users or to analyze or judge the scope of the content to be put on the WWW. The second reason is the disregard for, and inadequacy of methodology and guidelines for, designing and organizing information on the WWW. Finally, there is the SuperWebmaster syndrome. More and more frequently, a person publishing on the WWW, often referred to as a Webmaster, attempts to or is expected to do it all—HyperText Markup Language (HTML) mark-up, Common Gateway Interface (CGI) programming, graphic design, user interface (UI) design, information structuring, and Web site maintenance. Although possible, and understandably necessary for smaller Web sites, this situation often leads to a compromised design solution for medium- to large-sized Web sites.

OVERVIEW OF THE CHAPTER

In this chapter, a methodology for designing Web sites based on human factors and usability engineering principles is presented. This methodology focuses on having a clear understanding of user needs, that is, understanding who the users are, what tasks they need or want to perform, and what hardware and software environments they use. Emphasis is also placed on understanding the needs of content owners and authors in designing a successful Web site, an important consideration, as content owners and authors play a significant role in ensuring the ongoing usability of Web sites after their deployment. The importance of back-end processes that support interactivity and two-way communication at a Web site is also discussed. In this methodology, iterative testing of Web sites at various stages of design and development is considered key to meeting site objectives and to designing effective Web sites. The possibility of conducting remote usability evaluations to address a diverse range of user populations is also explored. Next, the critical nature of Web site maintenance is discussed. An attempt is made to convince readers that this stage determines success of the Web presence. Finally, the impact of future technologies such as WebTV and "narrowcasting" on design of Web sites is discussed.

METHODOLOGY FOR DESIGNING WEB SITES

The proposed methodology uses steps similar to those required for any system development: *planning, analysis, design and development, testing, implementation,* and *maintenance*. There are clear differences for designing Web sites, however, and they become evident in the description of each stage.

Planning

In the planning stage, the most important question is: "Why design a Web site?" Unless this question is answered satisfactorily, plans to create a Web site should not proceed. Once the goals of the Web site have been clearly identified and shared with the project team, the next step is to understand users and use environments. This understanding helps define the boundaries for the Web site design. Finally, it is important to identify the needs and requirements of content owners and authors. This last step is important, as owners are partly responsible for ensuring ongoing usability of the Web site after deployment.

Understanding the Goals of the Web Site

Identifying and establishing the goals of a Web site makes it easier to plan the overall organization and presentation of the site's information (Heller & Rivers, 1996; Sano, 1996). Without a clear purpose for a Web site, development very likely proceeds haphazardly, in a manner that Blundon (1996) referred to as "a process of random mutation." Take this, add that, shave a little off the sides.

The purpose of a Web site can be information sharing, communication and collaboration, interface to customers, electronic commerce, transactions and applications, or a combination of all these. In designing intranets (or internal Web sites), for example, a major goal may be to facilitate better communication and access to information for all employees in an organization (Sano, 1996). On the other hand, for online services such as America Online (http://www.aol.com), CompuServe (http://www.compuserve.com), and Microsoft Network (http://www.msn.com), the purpose of a Web site could be to get people engaged. The design would then be based on encouraging free play, fun, ease of use, and repeat use (Schuh, 1996).

Understanding Users and Use Environments

Once the goals of a Web site are established, the next step is to understand and capture user needs, usage characteristics of the content, and use environments. The way that a Web site is designed often depends on whether it is for external or internal users. For an external audience, a wide variety of client software, hardware, and network connections need to be accommodated. An internal audience allows a better understanding of hardware and software requirements. Furthermore, the business need drives the design of internal sites; an external site may need a slicker, more sophisticated presentation to bring users to it. For both internal and external users, it is important to know users' familiarity with computers, requirements for using and accessing information, and knowledge of and familiarity with the content and with Web technology. According to Odyssey Homefront Survey (1996), 40% of 2,000 users did not know which Web browser they used.

It is also important to know how Web enabled users are—that is, whether they have the necessary connectivity and client software to access the WWW. This information is particularly important for intranets: When all users do not

have access to the WWW, providing information solely on the WWW could seriously hurt a company's information sharing goals.

Finally, it helps to know what type of connectivity—bandwidth and network considerations—users have. In GVU's sixth survey, the most common connection speed reported was 28.8 Kbps (51.40 %), followed by 14.4 Kbps (19.69 %). Thus, with 70% of the user population accessing the WWW with slow-speed modems (and not high-speed connections such as ISDN or T1 lines, which can provide speeds of 128 Kbps and 1.54 Mbps respectively), designing a graphic-rich Web site could seriously affect its usability.

Identifying Content Owners and Authors and Understanding Their Needs. At the planning stage, the focus should not be exclusively on users of information. Before beginning to design, developers must identify content owners and authors and involve them in the process of Web site development. This process is necessary because the authors and owners are partly responsible for ensuring ongoing usability of a Web site after its deployment. If questions such as, "Who owns the information presented at the Web site?" "Who has responsibility for deciding what goes where?" and "Who decides what the Web pages should look like?" are answered before designing a Web site, especially for large-scale Web sites, many problems related to ownership and responsibility can be avoided at the maintenance stage. Also, it may be useful to get corporate public relations involved to make sure that appropriate branding or corporate image is conveyed by the Web site. When these considerations are overlooked, Web sites may have very short life spans.

Knowing who the authors are, how knowledgeable they are about computers and Web technology, and where they are located (at one location or at different locations) are all important to determine whether any training is necessary, and what approach is best for updating a Web server. There are three approaches for managing the content on the Web server: (a) centralized access, in which only an administrator has access; (b) distributed access, in which a team of authors manage a Web site without any access restrictions; and (c) limited access, in which authors have access only to certain areas of a server.

A growing number of Web sites are designed using object-oriented technologies in which individual pages are constructed on the fly. Some "intelligence" may also be built into such systems, so that the pages are customized—narrowcasted—for individual users. Training of authors and content owners becomes extremely important if such advanced technologies are used to develop and maintain a Web site.

Analysis

During the analysis stage, decisions are made about both content and process. Content-related issues refer to the structuring of information, the granularity of information at different levels of the structure, and the overall content navigation strategy. Process-related issues refer to how the content is maintained and how the interactive aspects of the Web sites are handled behind the scenes so that they are

transparent to users. The analysis stage, therefore, can be divided into two distinct substages: content analysis and process analysis.

Content Analysis

Content analysis requires identifying user tasks, understanding the information needs of potential users, and analyzing the content to be put on the WWW. It is important to define and analyze the content before spending a lot of time on page design. A designer often invests considerable time and effort in page layout and navigation design only to find later that the design does not work for the content or purpose of the Web site (Apple, 1996).

Identifying User Tasks. Content analysis must begin with identifying the tasks that a user may perform at a Web site. The tasks could be similar to the following (Heller & Rivers, 1996):

- Find a particular product.
- Learn about the newest products.
- Hear product reviews and customer testimonials.
- Order a product line.
- Learn about product features.
- Get product demonstrations.
- Contact organization employees.
- Obtain technical support.
- Provide feedback.
- Follow up an advertisement.

Knowing the tasks keeps design activity focused and free of extraneous information. The main objective during the design stages should be to provide mechanisms for performing tasks that align with the goals for the site and the information you have about users (Heller & Rivers, 1996). Later, users can be asked to perform the same tasks during usability evaluation of the Web site.

Determining What is Worth Publishing. The content of a Web site should be subjected to careful scrutiny of its purpose and objectives for users. If content does not have a purpose, such as to communicate, to educate, to inform, to entertain, it should probably be removed. The question that should be asked during content analysis is: What information is useful to users? The answer to this question helps determine inclusion and exclusion criteria for content during the maintenance stage as well (and it should be included in the Web site design style guide). Wilson (1995) considered this process to be deciding what constitutes information worth publishing. Often, developers want to put all available information on the WWW, without thinking about who uses it. If scrutiny is overlooked, users become frustrated by struggling through content to get to the information they need.

Organizing Information. The next important activity in content analysis is structuring or organizing the information. The organization of information should be based on users' mental models of the content structure. One caveat is in order: The site should be organized to support user tasks and goals of the site rather than internal management structure. Many Web sites are based on internal organization or management structure (Heller & Rivers, 1996), and Web site developers expect users to know this. If users must know an organization's internal structure to find information about its products and services, the Web site is unusable by a significant percentage of users. Such a site may be easier to maintain internally;

the roles and responsibilities can be clearly divided on the basis of department or division boundaries. But only users in the organization who understand the structure of the organization find such a site easy to use, and many employees in a medium-to-large-scale organization, would not know the corporate structure.

A common technique used to discover users' mental model of an information space is the card-sorting technique (Nielsen & Sano, 1995). In this method, users are given a set of cards in a random order and are asked to sort or organize them into different groups according to their similarity or their "belongingness" to a particular group. After the cards are sorted, users are asked to name each group. The group names (or category titles) are important as they help users form an expectation of what lies at the destination page. Category titles must appropriately describe information; otherwise users may not explore the Web site. It has been observed that users are reluctant to explore an application when they are unable to understand an icon or command (Rice, Farquhar, Piernot, & Gruber, 1996), but it is extremely difficult to think of category names that are understood easily by everyone. In a usability study on an intranet site, Vora (1996) found that users visited several categories to locate information in a search task, when it was not clear which category would have the desired information. Similar problems have been reported in the field of human–computer interaction for command names (Furnas, Landauer, Gomez, & Dumais, 1987). Therefore, it is often advisable to supplement category names with a brief description of the destination page.

In a usability study at US WEST Communications, the card-sorting task was extended by asking users to reorganize the cards in the groups on the basis of their importance in a group. The importance information can be useful in laying out navigational pages because it helps determine the order of information on a page.

Identifying Subwebs and Page Types. Once the organization of information is discovered, clusters of information may become obvious. Consideration may then be given to organizing these "clusters" of information into subwebs or what Nielsen (1996) referred to as sub-sites. Subwebs offer identical style characteristics and appear to users as a single site (Adelson, Rivenburgh, & Smith, 1996).

After the subwebs are identified, the next step is to identify the page types, which are identified on the basis of shared functionality or purpose. Some examples of page types are home pages for introducing each sub site, user log-in pages, feedback pages, index pages, personal home pages, search query–results pages, project home pages. Each known page type must be identified, along with its functionality, contents, controls, and cross-links (Sano, 1996).

In addition to identifying the page types, common headers, footers, and navigational controls should be determined; these may be common for all the pages at a Web site and for all pages in a subweb. For example, all the pages at a Web site may have links to the home page for the subsite, the Search page, the Feedback page, and the Help page; and all the pages may have copyright

information, the last update date, and the Universal Resource Locator (URL) of the page.

Having made these decisions for the page types and the Web site as a whole, developers should design templates for all the page types, with placeholders for each kind of information such as headers, footers, text, and graphics to ensure a consistent look and feel for the Web site. This not only facilitates Web site development, but also minimizes maintenance when new information is added.

Determining Update Process and Schedule. Finally, during the content analysis stage, it is important to determine who is responsible for providing fresh information for different areas of the site, how updates are communicated to users, and how frequently information is removed, archived, or updated.

Process Analysis

A Web site is not a static structure; it can be highly interactive (Wilson, 1995). Making a Web site interactive, especially to promote two-way communication, is often desirable. If a Web site is designed for interactivity, it is necessary to analyze the process(es) that support this interactivity. For example, supporting two-way communication at a Web site requires several process-related decisions, such as:

- Who receives the feedback?
- How is recipient notified (via E-mail or a Web front-end)?
- What type of action is taken on the feedback?
- Is the feedback categorized (comments, suggestions, questions, etc.)?
- Who responds to the feedback (are subject matter experts (SMEs) involved)?
- How are they notified that a response is needed (E-mail, Web front-end, or both)?
- How do they respond to the feedback (E-mail, on the Web)?
- Who makes the decisions to make changes?
- Who reviews (including legal review) and edits the response?
- How long should the response be?
- How are users (who provided feedback) acknowledged (E-mail, Web front-end, or both)?
- What is the turn around time (24 hrs, 2 days a week)?
- If users need a clarification of the response, how do they initiate it (threads or a separate feedback)?

These process decisions also affect how the Web UI is designed and what is promised to users. Without addressing such process-related issues, two-way communication initiatives may backfire.

Design and Development

In the design and development stage, information gathered in the earlier stages is translated into actual design. But before beginning to design, Web site designers should become familiar with Web technology such as Hypertext Markup Language (HTML), graphic compression schemes (e.g., GIF, JPEG), and common gateway interface (CGI) programs, and should understand their continually changing potential and limitations (Sano, 1996). In particular, designers need to understand the way graphics and multimedia work on the WWW. They need to understand technological issues such as compression schemes of graphic and multimedia formats (e.g., GIF, JPEG, QuickTime, WAV), access or response times, and cross-platform constraints (Weinman, 1996; Wilson, 1995). Designers must ensure that Web pages display well on a variety of computers with different color palettes and must take into account users' differences in network connection speeds (Wilson, 1995).

As with hypertext systems, the major problem facing Web designers is managing complexity: how not to overwhelm users with a vast amount of information. Understanding how Web users behave can be very useful in designing comprehensible and easy to use Web sites. In addition, several other design issues that need to be considered when designing a Web site, such as use of metaphors, design of individual Web pages, use of graphics and multimedia, and use of advanced technology.

Designing for Users: Understanding User Behaviors

To help designers, several design guidelines are available on the WWW:

- Ameritech Web Page User Interface and Design Guidelines (http://www.ameritech.com/news/testtown/library/standard/web_guidelines/introduction.html)
- Apple Web Design Guide (http://www.cybertech.apple.com/HI/web_design/intro.html)
- IBM Web Design Guidelines (http://www.ibm.com/ibm/hci/guidelines/web/ web_design.html)
- Sun's Guide to Web Style (http://www.sun.com/styleguide/)

With changing technology, however, these guidelines need revisions and designers must be informed about the changes so that they can implement better design solutions. To address this problem, on the basis of published guidelines and usability studies on the U S WEST intranet sites, I have identified 12 fundamental Web user behaviors and preferences that are likely to remain unchanged, despite changes in Web technology. Table 12.1 lists both the behaviors and examples of corresponding Web design guidelines. Web developers can use these behaviors and preferences to guide their design decisions when design guidelines offer no assistance.

TABLE 12.1
Examples of Design Guidelines for
Commonly Observed Web Reader Behaviors and Preferences

Reader Behaviors & Preferences	*Example Design Guidelines*
Users do not read the entire page. They scan the links first.	Use meaningful text in your links. Avoid linking the words Click Here.
Users do not like to read text online.	Allow for expansion of information using hypertext links. Make the text easier to read.
Users care about their time.	Avoid graphics that do not add value to Web pages. Indicate width and height in the IMG tag
Most users are casual users of your Web sites.	Use a simple structure not more than 3 levels deep. Indicate what's new on the Web site.
Users may be introduced to a Web site through pages other than the home page.	Create an identity for a Web site. Create a link to the home page.
Users do not want to spend unnecessary effort.	Avoid using banners that are wider than standard window size.
Users are interested in only a small part of the overall content at a time.	Keep navigation simple. Create a distinctive but consistent look for Web pages.
Users get lost.	Provide signature information. Allow users to return to the home page from every Web page.
Users appreciate pages that are pleasing.	Use eye-catching graphics, but not "in your face" graphics. Create a family of graphics.
Not all users have the same client software (Web browsers).	Design to support text-based browsers. Design for browsers at least one version lower than the most current version.
Not all users have the same type of hardware platform.	Use non-dithering color palette for designing graphics. Use banners no wider than 475 pixels so that Macintosh users do not have to resize their browser windows.
Users care about ease of use.	Do not use technical jargon. Do not require navigation more than 3 levels deep.

Use of Metaphors

Like software applications that use desktop or book metaphors, many Web sites are designed with metaphors to make the sites comprehensible. Metaphors are intended to help users understand the overall organization of the Web site and to help them locate information.

Use of metaphors should, however, be carefully considered when designing a Web site. First, unless a metaphor has the potential to be universally understood, its impact may be limited and confusing, at least to those unfamiliar with it. Another important factor is the likelihood of adding new *types* of information at a Web site. If there is a good chance that new information types

will be added at a Web site, metaphors may be restricting; additional information types may not fit into the existing metaphor, and the authors may have to stretch their imaginations to decide where to put the new content into the existing metaphorical organization. Users may also find it difficult to locate information. In some situations, it may be easy to accommodate new content by extending the metaphor. But, if a visual metaphor is used, it may be necessary to change the graphic to accommodate the extension of the metaphor.

This discussion is not intended to discourage use of metaphors or to suggest that metaphors should be abolished from Web site designs. The intent is simply to point out some limitations of metaphor use, especially for Web sites whose scope is likely to change often.

Design of Individual Pages

As most Web users scan a page rather than read it, design of individual pages becomes very important. In my experience, designing individual pages is the most time-consuming activity during development. Although a lot can be learned from visual design principles of layout, assistance is forthcoming from empirical research as well. Nygren (1996) identified two factors that affect search time on a Web page:

- *Number of items on a page* (search time increases with number of items).
- *Number of vertical groups or columns* (search time increases with number of groups). However, the number of groups must be traded off with the need for scrolling. The cost of scrolling in some situations may be higher than having multiple vertical columns. In such cases, using multiple vertical columns may be beneficial if it avoids scrolling.
- *Sorted versus unsorted groups and lists* (search increases for unsorted groups).
- *Horizontal groups* (search time increases rapidly with number of items in a horizontal group).
- *Number of features such as color, shade, slant, size, and space* (average search time increases with more features on a page).

Nygren (1996) has proposed a mathematical model to determine search times from the number of items in a page, the number of groups and sorted or unsorted lists, and the potential interactions between these factors. The model, however, needs to be validated with field studies and experiments. Until then, the influence of the listed factors should be considered when designing layouts of Web pages.

Use of Advanced Technology

Designers must balance the use of advanced technology against users' need to access the information and against constraints imposed by users' existing

hardware and software. When Web designers use sophisticated technologies, they also must decide whether a large number of viewers will be able to take advantage of them. From the information gathered during the planning stage about the use environment, it may become clear whether use of advanced technology is feasible. Even when feasible, advanced technology should be used only if it considerably improves usability and offers a richer experience to users.

Use of Graphics and Multimedia

There is no denying the importance of good visual design and attractive graphics on the WWW. When designing a Web site, developers need to consider graphic design aspects such as color, layout composition, and style. I have yet to see a user who despises Web sites that are visually appealing but that do not compromise interaction efficiency. Therefore, to ensure that interaction efficiency is not risked, designers need to understand the way graphics work on the WWW. As mentioned before, designers must know technological issues such as compression schemes of graphic and multimedia formats (e.g., GIF, JPEG, QuickTime, WAV), access or response times, and cross-platform constraints (Weinman, 1996; Wilson, 1995). Designers must ensure that Web pages display well on a variety of computers with different color palettes and must take into account differences in speed of access (Wilson, 1995).

Attractive graphics are not sufficient to make a Web site successful. The Web site of *The Chicago Sun Times* is a good example. As reported in Ragan's Webmaster Report (1996), when the Web site was released in October with outstanding graphics and slow downtimes, a consistent request from readers was to remove the graphics. The Web site was redesigned so that it was fast and accessible to users at any modem speed. A Web site design must keep a good balance between use of graphics, multimedia, and content to accommodate the needs of those who want fast and well structured quality information (content), and those who also want a pleasing and attractive design. As Wilson (1995) asserted: "The goal is to create cohesive systems; visual appearance is only one factor. Web designers who restrict their efforts only to how their Web pages look tap only a small part of the Web's potential" (p. 2).

If multimedia is used (and streaming technologies are not used), it is important to indicate format, file sizes, and need for any special helper application to view, hear, or interact with the media, next to the multimedia document. If possible, a designer may also indicate the time it takes to download using 14.4K or 28.8K modems for those who do not have a high-speed connection to the Internet.

Accessibility Related Design Issues

To reach a diverse user population, Web site design should be accessible. As this issue is discussed in chapter 7, it is not addressed here, but readers are advised to consider accessibility-related issues in their designs.

Designing for International Audiences

Web design changes are often necessary to address regional, linguistic, cultural, and writing system differences. If the objective is to reach international audiences, appropriate measures should be taken from the beginning. To design for an international audience, the following guidelines from Apple Web Design Guide (Apple Computer, Inc.; http://www.cybertech.apple.com/HI/web_design/intro.html) should be considered:

- Provide a way for users to choose the language.
- Know the region for which the Web site is developed.
- Choose the graphics wisely.
- Know the date.
- Consider language differences; do not use slang to avoid possible confusion misinterpretation, or mistranslation by translation engines.
- Avoid assumptions about text.
- Consider alignment differences.

Design Style Guides: Creating Consistent Web Sites

An agreed on style guide for designing Web pages, like OSF/Motif, Windows, and Apple Macintosh design style guides, does not exist. To create a clearly identifiable and predictable look and feel of a Web site, a Web design style guide should be developed for Web authors. The guidelines should address look and feel and navigational usability, as well as maintenance of Web sites. Specifically, the following information should be included in a design style guide:

- Description of each sub site
- Templates for each "page type"
  - header, footer and structure information
  - style sheet information
- Form layout examples and acknowledgment page
- Guidelines for adding content
  - identification of new content
  - placement in the existing organization
  - "newness" interval
  - file-naming conventions
- Guidelines for removing content
- Presentation guidelines
  - color schemes, layout templates, branding
- Decision aids
  - GIF versus JPEG, HTML versus PDF, etc.

Usability Testing

Iteratively testing a Web site at various stages of design and development is key to meeting the objectives of the site and to designing an effective site. The nature

of usability tests depends on how far the design has proceeded. During the early stages of development, paper prototypes and HTML prototypes can be tested quickly with "discount" usability engineering techniques (Nielsen & Sano, 1995). During later stages, more elaborate usability testing can be done by capturing time, error, and satisfaction data.

Discount and Opportunistic Usability Evaluation

Usability testing and evaluation need not be very formal, especially during the early stages of Web development. Although a lot can be learned from formal usability testing, it may be sufficient to follow a discount or an opportunistic evaluation strategy. Discount methods are deliberately informal and rely less on statistics and more on an interface engineer's ability to observe users and interpret results (Nielsen & Mack, 1994; Nielsen & Sano, 1995).

In addition to the discount approach, I propose the opportunistic approach—that is, test whenever and wherever possible, with co-workers, family, friends, or someone who stops by the office. Ask for comments on overall look and feel. This approach is based on my belief that any evaluation is better than no evaluation at all. The opportunistic approach should be complemented with the discount evaluation methods and formal usability testing whenever possible to evaluate a Web site completely.

Both discount and opportunistic methods are especially easy to use on the WWW. A major advantage of Web-based development is that a site can be tested online, while it is still in the design phase. Designing prototypes using HTML is fairly easy and can take place during the early stages of development to test the viability of designs and to evaluate the design assumptions.

Conducting Usability Testing and Evaluation

An easy way to conduct a usability test is to give users a set of tasks (including search tasks) to perform using a Web site prototype, and observe and take notes as they complete the task (Nielsen & Sano, 1995; Vora, 1996). Users may also be asked to think aloud or verbalize their actions and thoughts as they go along, so that their experiences can be captured. To capture quantitative data, it may be necessary to record their navigation in the background along with time stamps. But conclusions derived from time-stamped data should be considered carefully, as they may be quite unreliable. In a usability study done at U S WEST, although no time-based data were collected, considerable variability was found in the system response time depending on the time of day. Test sessions were conducted at 8:00 am, 10:00 am and 1:00 pm; the response time for the 8:00 am session was the most acceptable, and degradation in response time by 1:00 pm was very close to unacceptable.

The tasks or scenarios identified during the planning stage may be used as representative tasks for evaluating a Web site. It may be useful to include some free exploratory tasks as well, to capture casual users' behavior with the Web site. Thüring, Hanneman, and Haake (1995), and Stotts and Furuta (1991) have

suggested that any investigation of hypermedia cognition should include both free exploration and constrained (well-defined) tasks.

Finally, users may be debriefed about their experience and asked to fill out a satisfaction questionnaire to capture their interaction experience. Usability testing or evaluation of Web sites should focus on at least three dimensions: content, navigation, and presentation. A few examples of usability statements for evaluating Web sites are:

- It is obvious which category contains the information I need.
- Information loads very quickly at this Web site.
- Recently added information is clearly identified at this Web site.
- Web site X's Web pages are not cluttered.
- Icons used by Web site X are understandable.

Remote Usability Testing and Evaluation

Development on the Web offers an ideal opportunity to do usability evaluations remotely and to reach a wide cross-section of users. It is relatively easy to put the prototype on a test server and have users access it with their Web browsers, perform a representative set of tasks, and give their feedback to the site developers via E-mail or a form on the Web. Users may also fill out online satisfaction questionnaires. To follow what test users are doing, a simple solution is to use a telephone (preferably a speakerphone) to communicate with them and to use a product such as look@me (from Farallon Computing; http://www.farallon.com) to see the screens of remote users. As collaborative tools on the Web become more common, conducting remote evaluations may become an evaluation method of choice.

Scope of Usability Testing

A useful question in usability testing is: What should be the scope of usability testing and evaluation, especially when there are several hundred or more pages at a Web site? I believe that testing should be focused on home pages, navigational pages, template pages, navigation mechanisms and search mechanisms. Evaluation of individual pages can be done heuristically with a simple checklist-based approach or by a UI designer, for aspects such as consistency, layout, navigation, and orientation support.

A usability evaluator should be careful to distinguish between usability problems with a Web site and those with Web browsers. Current Web browsers do not offer the most usable interfaces. For example, studies by Tauscher and Greenberg (1996), and Jones and Cockburn (1996) clearly indicated the problems with stack-based history mechanisms in current Web browsers.

Off-Line Usability Testing

Several aspects of a Web site can be tested without involving users, especially those who attempt to validate presentation on different platforms and at different

connection speeds. Such off-line tests should be conducted throughout development to make sure that a site is accessible to a wide range of user populations. Specifically, the following should take place before a Web site is made available to users (Walters, 1996):

- Checking response time from the actual Web server.
- Testing the Web site with images turned off for navigability of the Web site.
- Testing the site with different Web browsers and platforms.
- Testing the site with a 256-color screen.
- Matching the titles of Web pages with actual content and anchor names that lead to these pages.
- Testing the site at lower access speeds (14.4K or 28.8K).
- Testing interactivity: password protection–registration requirements, search engines, database accesses, and so forth.
- Testing the Web server at expected peak usage.

Implementation

This step is probably the easiest in Web site development. In this stage, the content, whether developed locally or in a distributed manner, is moved to the Web server. Final testing of links and CGI programs are implemented to ensure that there are no dangling links and that all the CGI programs work as expected. For reasonably-sized Web sites, it helps to take advantage of Web site management software so that dangling links can be easily identified. Several current versions of site management software and Web servers offer such functionality.

With ongoing Web site development, before and after deployment, it is likely that dangling links exist and that some users encounter them. In such cases, users receive a File Not Found error. To minimize the impact of such situations, Strehlo (1996b) recommends developing a CGI program that overrides a generic File Not Found error message and shows users a predesigned Web page. This page provides a small description of the problem and asks users to fill out a form with information about what browser they are using, the platform they are on, and any details about how the error occurred. With this information, the Web development staff can identify and fix the problem and notify users via E-mail with the proper address of the information.

Another useful tool for Web site managers who are expecting millions of hits a day is a load-testing tool to check Web server performance with different levels of traffic to their Web site. Site managers can either rely on published Web server performance reports or use load-testing tools such SiteTest (from Mercury Interactive; http://www.merc-int.com), to generate virtual hits to test site performance. Site managers can also use the results to diagnose and solve production problems before deploying the Web site for general use.

Maintenance

Web site design is "a process and not an event" (Kawasaki, 1996), and maintenance a critical factor in having a successful Web presence. Maintenance is easier said than done. The challenge is to provide ever-changing, useful content while maintaining the integrity of the Web site. Maintenance also requires that Web site usage trends be analyzed and changes be made to accommodate users' needs. Finally, to keep the Web site interesting and useful, Web site managers must periodically identify, evaluate, and incorporate new technologies as they become feasible.

Updating the Web Site Content

An important task in Web site maintenance is updating the content and managing interactivity. If planning and analysis are done carefully, maintenance need not be a difficult chore. For instance, if the processes established during the planning stage with the content owners are followed and the design style guides are adhered to, ongoing content delivery is easy and consistent. Furthermore, if process owners are identified for supporting interactivity and responding to user queries and feedback, maintaining fluid two-way communications should be straightforward.

Maintenance also includes integrating additional content into a Web site and letting users know of the new content. Integrating content may seem easy for a small-sized Web site, with just one or two people managing the content. But, for a medium- to large-sized Web site, it may be useful to create a staging area on the server, which allows everyone involved to see what the new information looks like, to ensure a look consistent with the existing site and to approve content if necessary (Strehlo, 1996a).

Checking Web Site Integrity

A well-known maintenance problem is identifying and removing dangling links. As Web pages are moved or their names changed, the links that refer to them must be updated to reflect the change. Some Web servers support Web site maintenance by showing the entire site graphically and indicating missing page links; others offer functionality where the link-integrity checking feature automatically updates links as pages are uploaded. Furthermore, if the moved Web pages are high on users' revisit list, users must be notified as well, by putting a Web page indicating the move and a link to the new location and asking users to update their bookmarks.

Monitoring Web Usage Trends

Monitoring and analyzing Web usage information is another important activity after a Web site has been implemented. Web site managers must monitor traffic, performance, and usage patterns; manage the application, configuration, resources, and content; and perform load and quality testing, if their Web site supports a production environment. Several available software applications (e.g., net.Analysis Desktop by net.Genesis; http://www.netgen.com) analyze log files

created by Web servers and automatically produce reports that include statistical information, as well as graphs showing user activity by interest level in specific pages or services, user origin (local, national, or international), Web browser types used, and connection speed. These applications also provide reports that show trends in server usage, busiest times, and number of server requests during peak times. This information allows site managers to isolate problems in download times, server response, and visitor connections.

Monitoring Web usage helps authors determine what users find interesting and useful. On the basis of usage patterns, resources can be allocated strategically, but usage trends may not indicate problems in navigation. For example, if a particular area of the Web site is not explored at all, it may mean either that the users do not consider it important or that users do not know of its existence (or that it is difficult to get to!).

Evaluating and Implementing Newer Technologies

The maintenance stage also includes considering newer technologies such as JAVA, ActiveX, JavaScript, and VBScript, and using them where necessary to provide a richer interaction experience for users. This process is both important and essential as there will very likely be a common interface for both information-based hypertext documents and clientñserver type Web-based applications in the future, and the Web may become a unified front-end in a distributed computing environment.

THE FUTURE OF WEB SITE DESIGN

Although the previously discussed design activities will remain the same in the future, considerations in each activity may very likely change. The current trend is to customize the presentation and content for individual users, and to use WebTV (http://www.webtv.com) as a potential delivery channel. The impacts of these trends on Web site design are briefly discussed next.

Narrowcasting or "Personalizing" the Web Content

A recent trend is to personalize Web sites that users visit on a regular basis. Rather than providing information as a static site, which appears the same to all users, users are increasingly offered the option of customizing the content to meet their unique needs and interests. Delivering customized content requires knowledge of what users want. This knowledge can be obtained by creating a user profile in several ways. Content providers can solicit users to identify their interests and can use this information to customize the content for them. Or, they can, with users' permission, build a database of their interactions and extract a profile from the gathered information. Whatever the approach, future Web sites may have to consider gathering user profiles and matching them to different presentation styles very early in their design process.

Designing for the WebTV

If WebTV becomes popular, Web developers and designers must radically change their design approach to create Web sites that are optimized for the television medium. Changes are necessary, as WebTV users are different; their interaction style is different; their environment is different from those of Web users.

WebTV users are likely to expect the kind of interactivity they get when changing channels on their TV; they may not want to wait for long downloads to view a Web page. They may very likely use a remote control to interact with the Web and may have a strong dislike for scrolling. And unless they can get to the information quickly, they may simply abandon interacting with the Web. Furthermore, not all users will have a keyboard, and forms that require considerable typing may become unusable; responses must be simple, such as Yes and No, or Agree and Disagree.

Because a television cannot offer the same richness in resolution and colors as a PC monitor, designers must be very careful in designing Web pages. Many Web designers must totally rethink their design strategies. All is not lost, however. A few years ago, studies were conducted for interactive television (ITV). I am certain that much can be learned from these findings and observations to make sure that usability is not compromised for WebTV users.

SUMMARY

Although Web technology makes evolution, enhancement or changes to content relatively easy, Web site design cannot proceed in a haphazard manner. Web site developers must ask themselves whether they are putting up a Web site simply to have a Web presence or to deliver content and interactivity that is educational, informational, or entertaining, that helps improve productivity or lets users do something useful (Milano, 1996). The more thought, planning, and attention given to user needs during development and the more effort made to become familiar with benefits and limitations of the Web medium, the more likely that the result is an enriching, successful Web site (Sano, 1996). To address user needs, Web designers should analyze how users can use the information; ensure that users' questions are answered throughout a Web site; develop orientation and navigation to support users; and create standards that are consistent in function and screen design.

In addition to understanding user needs, usable Web site design requires understanding the needs of Web authors and content owners. Addressing the needs of authors ensures that users are consistently provided with updated content. This process can be facilitated by deciding on the model to update the content on a Web server, creating a Web design style guide, and clarifying roles and responsibilities. Special attention should be given to back-end processes to support two-way communication with users.

In the design and development stage, the information gathered in the earlier stages is translated into actual design. As about 70% of Web user

population interacts with the WWW through slow-speed modems (GVU, 1996), designers should carefully balance the urge to create visually appealing sites with users' need for efficient interaction. Trade-offs are required when considering use of multimedia and advanced technologies, as well. It is important to address accessibility issues for populations with special needs and for international users. To ensure usability, design assumptions and decisions should be tested with representative users early and often. The relative simplicity of Web technology makes this testing possible. The WWW also makes it possible to conduct remote evaluations to test with a distributed population of users.

Maintenance is a critical stage in Web site development. The challenge is to provide current and updated content while maintaining the integrity of a Web site. Maintenance also requires that Web site usage trends be analyzed and changes be made to accommodate users' needs. Finally, as technology facilitates greater design expression, latitude, and interactivity, user expectations increase. Web site managers should, therefore, keep up with the advances in technology and implement them to improve usability of a Web site and to offer a richer interaction environment for users.

Undoubtedly, Web technology is here to stay. It is not clear, however, whether browsing will remain the common way of interacting with the Web and whether desktop PCs will continue to be the platform of choice. Possibly, narrowcasting or customizing content and presentation based on specific user characteristics will become the norm. Web site designers will then need to invest significant time and resources early and continuously during Web development to determine user profiles and match appropriate presentation methods to them. Finally, WebTV is a strong candidate for the future platform of choice. If WebTV does become a popular delivery medium, Web designers must rethink their design strategy to accommodate the needs and expectations of totally different users.

REFERENCES

Adelson, S. J., Rivenburgh, R., & Smith, S. W. (1996). *Enforcing closure of subwebs and expansive Web sites.* Poster presentation at Hypertext '96 Conference, Washington, DC.

Apple Computer Inc., (1996). Apple Web Design Guide. http://www.cybertech.apple.com /HI/web_design /intro.html

Blundon, W. (1996). So you want to build an intranet? *JavaWorld.* http://www.javaworld.com/javaworld/jw-12-1996/jw-12-blundon.html

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

Graphic, Visualization, and Usability Center, Georgia Institute of Technology (GVU). (1996). GVU's 6th Web User Survey. http://www.cc.gatech.edu/gvu/user_surveys /survey-10-1996/

Heller, H., & Rivers, D. (1996). So you wanna design for the Web. *Interactions, 3*(2), 19–23.

Jones. S. R. A., & Cockburn, A. J. G. (1996). A study of navigational support provided by two World Wide Web browsing applications. In *Hypertext '96 Conference Proceedings*, 161–169.

Kawasaki, G. (1996, May 6). Four things to ask your Webmaster. *Forbes,* 128.

Milano, D. (1996). Myriad confrontations. *InterActivity, 2*(13), 4.

Nielsen, J. (1996, September). The rise of the sub-Site. *Alertbox*. http://www.sun.com /columns/alertbox/9609.html

Nielsen, J., & Mack, R. L. (1994). *Usability inspection methods*. New York: Wiley.

Nielsen, J., & Sano, D. (1995). SunWeb: User interface design for Sun Microsystems' internal Web. http://www.sun.com/technology-research/sun.design/sunweb.html

Nygren, E. (1996, October). *"Between the clicks": Skilled users scanning of pages*. Paper presented at Designing for the Web: Empirical Studies. Microsoft Corporation, Redmond, WA.

Odyssey Homefront Survey (1996, September 30). Browser war? What browser war? *Information Week,* 12.

Rice, J., Farquhar, A., Piernot, P., & Gruber, T. (1996). Using the Web instead of a window system. In *CHI '96 Conference Proceedings*, 103–110.

Sano, D. (1996). *Designing large-scale websites: A visual design methodology*. New York: Wiley.

Schuh, E. (1996, October). *Developing an Internet online service: Goals, methodologies, and solutions*. Paper presented at the Designing for the Web: Empirical Studies conference, Microsoft Corporation, Redmond, WA.

Stotts, P. D., & Furuta, R. (1991). Hypertext 2000: Databases or documents? *Electronic Publishing, 4*(2), 119–121.

Strehlo, K. (1996a). Three important Web tools: Process, process, process. *Infoworld, 18* (45) http://www.infoworld.com

Strehlo, K. (1996b). Two best practices in easy Web-site development. *Infoworld, 18*(42). http://www.infoworld.com

Tauscher, L., & Greenberg, S. (1996, October). *Design guidelines for effective WWW history mechanisms*. Paper presented at the Designing for the Web: Empirical Studies conference, Microsoft Corporation, Redmond, WA.

Thüring, M., Hanneman, J., & Haake, J. M. (1995). Hypermedia and cognition: Designing for comprehension. *Communications of the ACM, 38*(5), 57–66.

Vora, P. R. (1996, October). *Usability evaluation of the U S WEST Global Village*. Paper presented at Designing for the Web: Empirical Studies, Microsoft Corporation, Redmond, WA.

Walters, C. (1996). *Web concept and design*. Indianapolis, IN: New Riders.

Weinman, L. (1996*). <designing web graphics>: How to prepare images and media for the Web*. Indianapolis, IN: New Riders.

Wilson, S. (1995*). World Wide Web design guide: Learn to design professional Web pages*. Indianapolis, IN: Hayden.

Part IV: WEB RESEARCH AND DEVELOPMENT

13 Web User Interface Development at Oracle Corporation

Anna M. Wichansky
George Hackman, Jr.
Oracle Corporation

The World Wide Web (WWW) was created as a multimedia communications network to help scientists collaborate, free of geographical and time constraints. It has developed as a means for anyone with a computer and a telephone line to do work or find information by browsing through its graphical pages and following hyperlinks. The software industry is embarking upon the third stage of the Web's evolution as a development platform for a wide variety of applications. These applications will be used on thin client machines supported by powerful servers.

As the world's premier server software vendor, Oracle Corporation has a vested interest in developing server capabilities to benefit enterprise customers. Oracle has taken on this challenge by trying to "Web-enable" many of its products, and by leading the industry in the development of a thin desktop client called the Network Computer (NC)™.

To truly enable the Information Age, Oracle must develop software that is easy to learn and use by a wide variety of people with differing degrees of computer experience. This chapter describes how the Usability and Interface Design Group at Oracle developed and tested user interfaces (UIs) for three Web products: InterOffice™, a collaborative work software suite for end-users; the NC, a thin client for Web browsing and applications; and Enterprise Manager™, a database systems management tool. We share the common threads of collective wisdom that we have gleaned about usability in Web design to date from this work. We also provide a preview of our future efforts in developing Web applications.

ORACLE'S VISION FOR THE WORLD WIDE WEB

Oracle Corporation is the leading worldwide vendor of database software and information management services. Its products include the Oracle 7™ relational database, a complete suite of developer tools, and application products for accounting, human resources, and manufacturing. The company employs over 30,000 software professionals worldwide with gross revenues of approximately $4.2 billion in fiscal 1996.

Oracle's President and CEO, Lawrence J. Ellison, has articulated a distinct mission for the company related to the rise of network computing and the Information Superhighway. He has envisioned the decline of expensive, complicated personal computers ("fat clients"), which currently interact with servers, typically mainframe or high-end workstations running databases. The fat-client model of computing requires a big investment in time and money on every desktop. Personal Computers (PCs) run platform-specific operating systems; applications must be ported to each platform; and software must be administered at every user's seat. Ellison (1996) has proposed a "thin-client" NC model, in which desktop machines require little computing power and are easy to install and use. The NC has little or no local storage; users access software applications running on Web servers. This saves the enterprise customer the trouble of buying, installing, and upgrading thousands of client machines, since applications need only be maintained on the server to provide access to all clients. It is estimated that this approach may save enterprise customers $5000 to 9500 per year per client seat (Hof, 1996). This could make access to NCs virtually universal for schools and other customers where today's PC costs are prohibitive.

To support these changes, Oracle has proposed a Network Computing Architecture (NCA). This is a cross-platform infrastructure for developing and deploying object-based, network-centric applications, based upon open standards such as Common Object Request Broker Architecture (CORBA 2.0) and HyperText Transfer Protocol/HyperText Markup Language (HTTP/HTML). Programming objects would plug into a software bus, which could accommodate any standards-compliant client, application server, or database server. NCA would provide a means of integrating new and pre-existing software components. It would also provide a common user interface language, HTML/Java (Java by Javasoft, http://java.sun.com) across many applications, which could facilitate ease of use (Oracle, 1996). NCA has been endorsed by many hardware and software vendors that anticipate major business opportunities resulting from the paradigm shift toward network computing.

Role of Oracle's Usability and Interface Design Group

Oracle's Usability and Interface Design Group consists of 30 visual and interaction designers, usability engineers, and information designers who develop UIs for Oracle tools, server products and applications. Our mission is to make Oracle products easy to learn and use. We follow a user-centered design process which includes: customer visits, requirements generation, prototyping, scenario-based focus groups, UI design specification, heuristic evaluation, iterative usability testing and design, and competitive usability testing following product release. A complete description of our design cycle is available in Wichansky, Friedland, Slater, and Dirik (1996). The department maintains six usability labs, including an Internet Classroom and Internet Living Room, and a database of 1,400 participant volunteers. All efforts are directed toward improving usability through UI design and development of learning and information products such as computer-based training. Usability is operationalized in terms of user performance

and acceptance metrics which are quantified in each test iteration and compared with previous products to measure the effectiveness of our efforts.

WEB DEVELOPMENT EFFORTS

Oracle has developed several products which contain a WWW application component.

InterOffice™

InterOffice™ is a suite of groupware applications which supports electronic mail (e-mail), scheduling, directory services, and document management capabilities. It evolved from two earlier products, Oracle Office™ and Oracle Library™. InterOffice consists of a MAPI™-compliant back-end application which allows installation of document and e-mail repositories on a server, and Windows 95 and Web clients. Early in the project we agrced with developers that the two clients should be as similar as possible in terms of functionality and UI design, to promote positive transfer of training between them. This proved to be a great challenge, as we were working with early versions of commercial Web browsers which provided little flexibility in terms of frames, panes, and other windowing features we wanted to include. Early designs of the Web client followed a visual metaphor similar to the Windows 95 client, in which the major product functionality was divided among tabbed, hyperlinked notebook pages. A major obstacle arose in the lack of availability of a word processing application which could be accessed inside the Web. (At press time, Oracle, Netscape and other vendors have them in development. Applix [http://www.applix.com] has a word processor available that uses Java: Applix™ Anywhere Office™.) In general, however, the two interfaces were remarkably similar, as shown in FIG. 13.1.

We used a technique in which we enacted usage scenarios to demonstrate the prototype interfaces to focus groups of target market users. A scenario was developed depicting the interaction between an East Coast soft drink manufacturing executive and a West Coast marketing manager. The scenario included: e-mail exchanges with attachments (including video files), co-authoring of documents, remote download of e-mail and documents to portable computers, scheduling video teleconferences across time zones, and access to personal and corporate directories. The UI designer and the usability engineer enacted these roles, one on a portable PC with prototype software depicting a Web UI, and the other on a standard PC with a Windows 95 interface prototype. The major positive feedback from these groups was how much they liked the Web interface, how attractive the visual design appeared to be next to the Windows 95 client, and how easy it would be to sell the Web client into their organizations. (Note, they had not actually used the Web client, just watched a demonstration of it.)

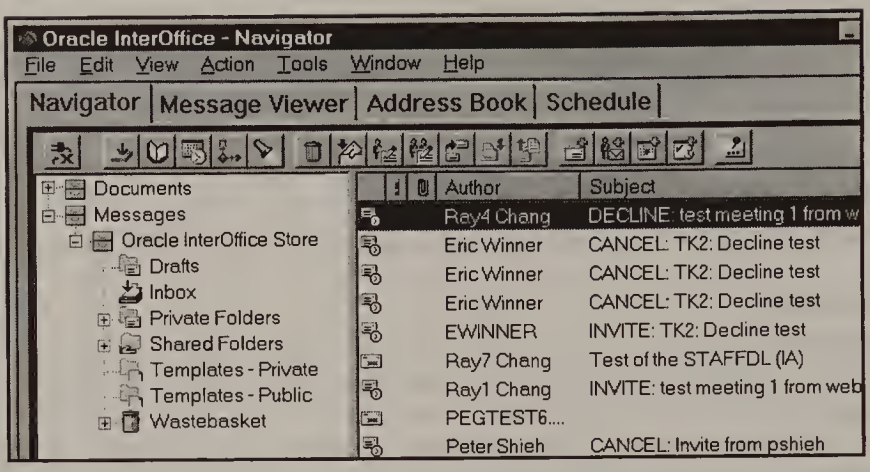


FIG. 13.1a. Oracle InterOffice Windows95 client.

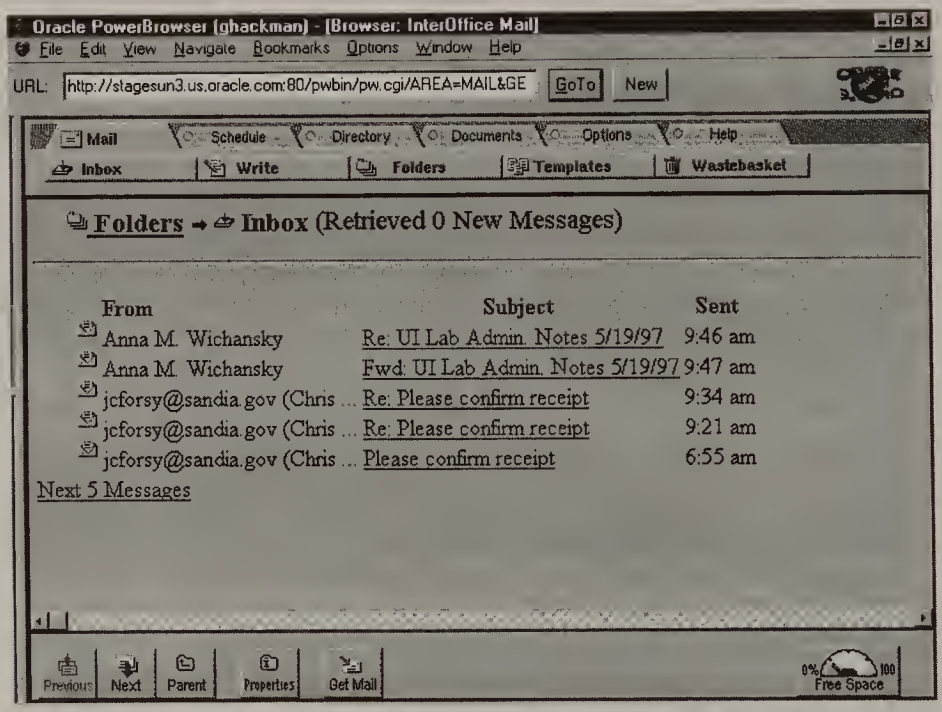


FIG. 13.1b. Oracle InterOffice web client.

After the focus groups, Windows 95 and Web client development continued. Both teams had design consultation from the usability group while developing the products. A UI designer worked closely with the Web development team in the design and layout of all Web UI elements. In fact, most of the actual layout and construction of the screens was done by the designer with the development team plugging in the functionality behind the graphical elements

in the screen. In several months both Web and Windows 95 products were ready for usability testing.

The first usability test of the Web and Windows 95 InterOffice clients commenced in the spring of 1996. At this point the two products had significant portions of their functionality implemented, but were not feature- or function-complete. The Windows 95 client had more of its core functionality implemented, including all e-mail functions, almost all scheduling functions, and some document management functions. The Web client had most of its e-mail functionality implemented, some document management, most of the corporate directory, and a few scheduling features.

The two usability tests were similar in design and execution. Each study employed eight business computer users (not computer programmers) who had experience with e-mail, word processing and the client platform (either Windows 95 or the Web). In both tests, users completed typical groupware tasks such as reading, sending and replying to e-mail, scheduling appointments, sharing documents and looking up users in a corporate directory. Since much of the scheduling functionality was not complete for the Web client, users were asked to tell us which pushbuttons or commands they would use to accomplish the scheduling tasks. Users followed basically the same testing protocol with a pre-test interview, a set of tasks they performed independently, and a post-test interview in both tests.

Both tests were successful in generating feedback on the product's UI, but the Web product performed better in the usability test. Users' satisfaction was higher for the Web product on three of four composite usability satisfaction measures (the two products were the same on the fourth measure). Users made fewer errors using the Web product, and they required less time and fewer assists to complete the tasks.

The types of problems encountered in the usability test for the two products reflected the design issues facing the two products. The most common problem encountered with the Windows 95 client was window management. In the Windows 95 client, each product function was developed as a separate OLE custom control that was launched when a user opened a schedule, address book or other function. The product allowed a user to open multiple OLE custom control windows simultaneously, as in FIG. 13.2. Users often had six or seven different windows open at one time; e.g. two e-mail message windows, three schedulers and two address books. Although this feature might be appreciated by power users, it confused novices who did not realize that they already had an open version of a particular window in the task bar at the bottom of the screen.

The primary problem that users encountered with the Web client stemmed from the lack of functionality within HTML. Users with limited Web experience expected most of the functionality of the Windows 95 interface in the Web, such as double-click activation, drag and drop, right click to pop up a context menu. The first few minutes of testing a user with little Web experience was always an adventure. Users would double-click everywhere in the UI and would not wait for

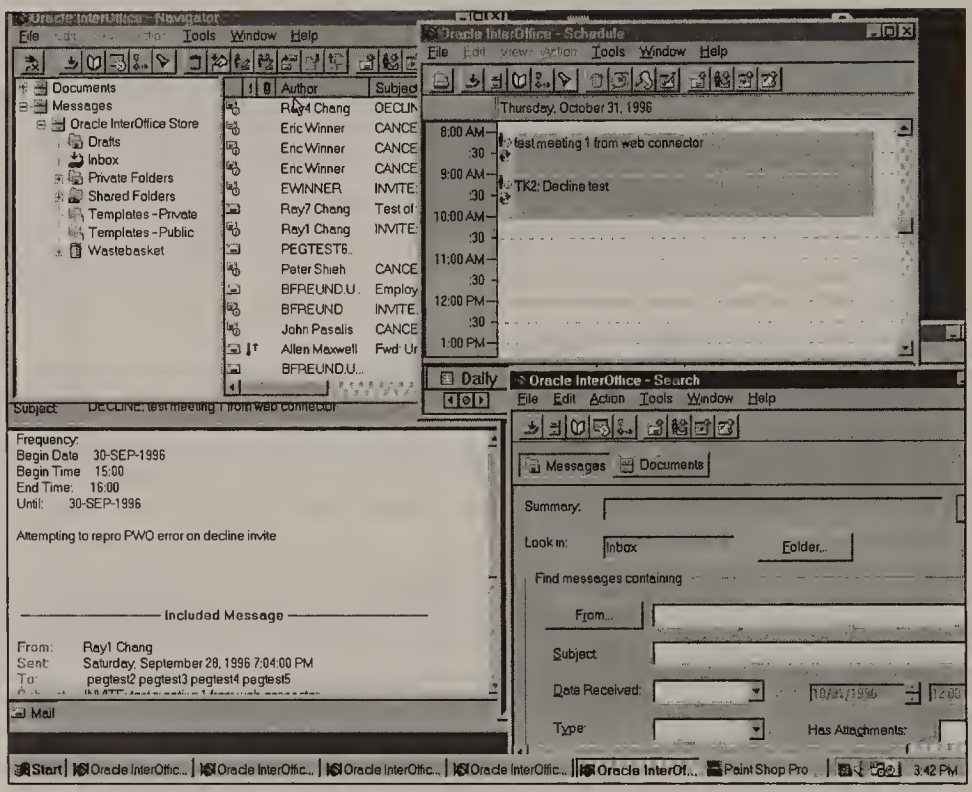


FIG. 13.2. Oracle InterOffice windows95 client showing multiple OLE custom control windows open.

the results to come back. After a few minutes of double clicking everywhere, the user would start to wonder "why the Web is so slow."

Less experienced users confused functionality accessed through the Web browser with functionality in the application itself. The limitations of HTML as a development platform caused developers to depend on the browser to handle some end-user functions, such as printing and searching for e-mail, that would normally be in the application itself. However, users did not think to look in the Web browser menu to find "print" or "search". Sometimes users ran into the opposite problem by looking in the browser for functionality that was in the Web portion of the application. For example, when asked to look up an employee phone number, some users would click on the "People" button in the Netscape Navigator™ (Netscape: http://www.netscape.com) instead of looking in the directory function in the application.

Users experienced with Web surfing tended to use the back and forward buttons frequently to navigate between pages within a given site. However, if users clicked the back and forward buttons within the InterOffice Web application, they would often return to out-of-date pages. For instance, some users would open a message from a list of e-mail, delete that message, and then use the

back button to navigate back to the e-mail list. Once users got back to the list, they were surprised to see that the e-mail they had just deleted was still there.

One unexpected usability problem encountered during the Web test involved data loss. When composing e-mail, users would often click the "Directory" tab to get the addressee's user name. When the users then clicked back on the e-mail tab to add the newly found address to the "To:" or "CC:" list, they would be surprised to find that the e-mail they were working on had disappeared and they were returned to the inbox. The cause of this problem was an HTML limitation related to the use of subframes in the UI. In reality, the e-mail was still there, but on a different page. The only way users could get back to the previous screen without losing their data was by using the browser's "Back" button. However, all the users in the test thought the e-mail had "disappeared." There was no complete fix at the time the product was first released, because of the limitations of HTML. We could only patch the problem with the error message shown in FIG. 13.3. Future releases of the product running in browsers with frame level caching should not have this problem.

Several months later a second round of usability testing was conducted on the feature- and function-complete versions of the two products. The task set for both products included a retest of the basic e-mail functions combined with a new test of functions that were not ready for the first usability test. For the Web product this included scheduling, creating aliases and mailing lists, and document management tasks. For the Windows 95 product this included features such as document management, e-mail filtering, rules, templates, and preferences.

For this test, users rated the Web client higher than the Windows 95 client on all the user satisfaction composites, but particularly the "ease of use" and "clear and understandable" measures. These results were probably due to the UI complexity of certain features of the Windows 95 client such as rules, document management, and preferences, many of which were not implemented in the Web version. With the Web client some of the same problems seen in the first test reappeared. This finding was not surprising, since users' expectations that the product should behave like Windows and confusion between interaction features in the browser versus the application were difficult to fix. Most new issues discovered were minor UI glitches that were easy to fix, such as mislabeled buttons or fields and lack of clear feedback.

The Web product shipped several months later with the confidence that as the Web matured the usability problems inherent to the platform could be fixed in

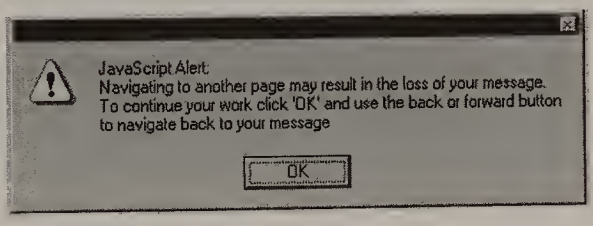


FIG. 13.3. Error message for Oracle InterOffice web client.

future releases. In fact, the Java-enhanced version of this software currently under development at Oracle has few of the UI problems we discovered in the Web version described above.

NC

The NC is a low-cost information appliance that enables the user to access the Internet via a telephone line or network connection. It can be used with a television (NTSC or PAL) display, remote control, and optional infrared (IR) keyboard, or a computer (VGA) monitor with standard cabled mouse and keyboard. It comes with its own Internet browser software and built-in word processor. It uses a credit card-like device as a means of login and authentication. This SmartCard™ identifies the user and allows his or her preferences and privileges to be recognized on any NC device.

Early in the project, we developed usage scenarios for the NC. In a home scenario, a user surfed the Web shopping for a sailboat. This scenario required: login, authentication, setting up bookmarks, entering Web site addresses (uniform resource locators [URLs]), selecting bookmarks, accessing Web sites that included databases with a bank account and small interactive financial programs, reading and sending e-mail, and logout. In a business scenario derived from visits to large Oracle customers, a biotechnology scientist accessed databases containing drug test results, used a spreadsheet application on the Web to analyze the data, and communicated the findings to remote colleagues and government agencies. We also developed multiple visual and interaction designs that would support these kinds of functionality. These designs were drawn up as storyboards and later rendered as partially interactive screen prototypes. These exercises were instrumental in the development of a complete UI specification for the NC, which involved the basic architecture for the visual design and interaction design, including the toolbar, icons, menus, help, remote control features and functions, and keyboard commands. The specifications will be supported in NC versions 2.0 and later, which include Java productivity applications.

There were many major challenges in working on the NC UI. First, the visual design had to work on both televisions and VGA displays. We consulted an experienced Hollywood video production designer to develop a strategy for designing the NTSC (i.e., television display) interface, which used different fonts, different colors, and a simpler interactive model than did the VGA design. Second, the major functionality had to be accessible using a remote control, or the standard computer mouse and keyboard. Typing with a remote control required screen interface design support. Finally, the greatest challenge was not knowing what users would expect to do with the NC. It was a new category of device, unavailable for data collection with users during development. We applied our collective experience from end-user computing, game devices, and interactive television design to the problem of anticipating NC users' potential wants, needs, and usability problems.

Once pre-release hardware became available, a series of NC usability tests was conducted in the Oracle Internet Living Room Lab. This facility was capable

of supporting both NTSC and VGA configurations of the NC simultaneously. Twenty users participated in the first test. Half the users had Internet experience and half did not. Half were home computer users and half were business users. The home users sat in an arm chair and worked with the NC on a 76.2 centimeter (30 inch) diagonal screen television, with a mouse and keyboard cabled to the device, as the remote control was not yet available. The computer users worked with the device at a desk with a mouse and keyboard and a VGA display. Both configurations were connected to the Internet via LAN; the modem was not yet available.

Users logged in, explored and performed a series of tasks typing in URLs, setting up bookmarks, reading Web pages, and word processing. The results showed that although Web users had reasonable expectations of the NC based upon their Internet experience with browsers running on PCs, the non-Web users were sometimes disappointed. These feelings were more related to the Web itself than to the NC. Users wanted to double click rather than single click hyperlinks. They thought the pages loaded too slowly (even on Oracle's LAN) and expected the fast local response times they got from PCs. Users had difficulty performing their tasks on the Web, and often got lost or disoriented when pages changed as a result of selecting hyperlinks. They also had a hard time dealing with long pages, especially when these had interactive controls on the bottom. Often they neglected to scroll long pages and accepted whatever landed on the screen as the complete Web page. They had difficulty searching. They sometimes expressed the feeling, "Is that all there is?", after having read the media hype about the Web. Web users' expectations were more realistic; in fact, some appeared to be already jaded about some of the Web's limitations such as response time.

Enterprise Manager and Applets

Oracle's focus on Web-enabling its products extends beyond the realm of end-user software to system and database administration products as well as database tools. An early effort was undertaken to Web-enable Oracle's Workgroup Server™ for database administration. The usability team became involved in this effort as a result of a desire for consistency across Oracle's database administrative products. At the time, a product called Enterprise Manager™ was under development and would become the mainstream Oracle database administration tool.

The UI group and executives at Oracle thought that the Web version of Enterprise Manager under development, ought to have a consistent UI with Enterprise Manager running on Windows NT™ or Windows 95. Because of the late stage of the development cycle for the first release of the Web product, we decided to focus on consistent functional groupings and terminology between the two products and a common set of graphical elements. Functions for the Web administrative product were regrouped according to the applet model employed by Enterprise Manager. Icons were redesigned to provide a consistent look and feel across the two products.

For the second release of the Web version of Enterprise Manager, the UI group was able to work with the development team in the initial design stages of the product. This early involvement, coupled with the more robust functionality of Java, allowed us to design a Web version of Enterprise Manager that was more consistent with the Windows version of the product.

In Enterprise Manager there was a console which consisted of four panes and a set of launchable applets (FIG. 13.4). The console was a power user tool that provided an overview of the network nodes, databases, and users throughout the system. It also allowed the capability of scheduling jobs and monitoring events triggered by the database. The applets were task-specific mini-programs that allowed users to perform specific groups of tasks, such as modifying database security and performing database backups

To replicate this functional grouping on the Web, the UI employed a preliminary log-in console with a set of Web hot buttons that users clicked to switch to the applet (FIG. 13.5). Like the Windows console, the Web console had four panes: the navigator, the map, jobs, and events. To account for screen space limitations of running in a browser, however, only the navigator pane and one other pane could be displayed at the same time. Like the Windows console, the Web console allowed the user to monitor the status of databases, events and jobs, but for the product's initial release, the Web console was read-only. To make changes to the monitored databases, users had to switch to one of the database applets by using the applet toolbar. To create or alter jobs or events they had to use the Windows version of the product. Future releases of Web Enterprise Manager will allow use of all the capabilities of the Windows client from any desktop with a compatible Web browser.

The released version of the Web administration tool provided the database administrator with an opportunity to do basic database management on any client running a Web Browser, without having to install and configure the Windows NT client. For the UI group, this product provided an opportunity to begin understanding the issues behind designing Web administrative applications. This knowledge has subsequently been instantiated as a specification for designing

Common Threads

Across our collective UI research efforts at Oracle, we have observed many common threads on user experiences and expectations for the WWW.

Great Expectations. The media hyperbole about the WWW has raised users' expectations about its functionality and ease of use. This has created a halo effect observed in the inflated subjective ratings of products with Web interfaces relative to what would be expected from user performance or traditional platform products with comparable functionality. It may also reflect the boredom that users experience with current unified platform environments.database administration products for the Web and is driving the design of more advanced Web administrative tools.

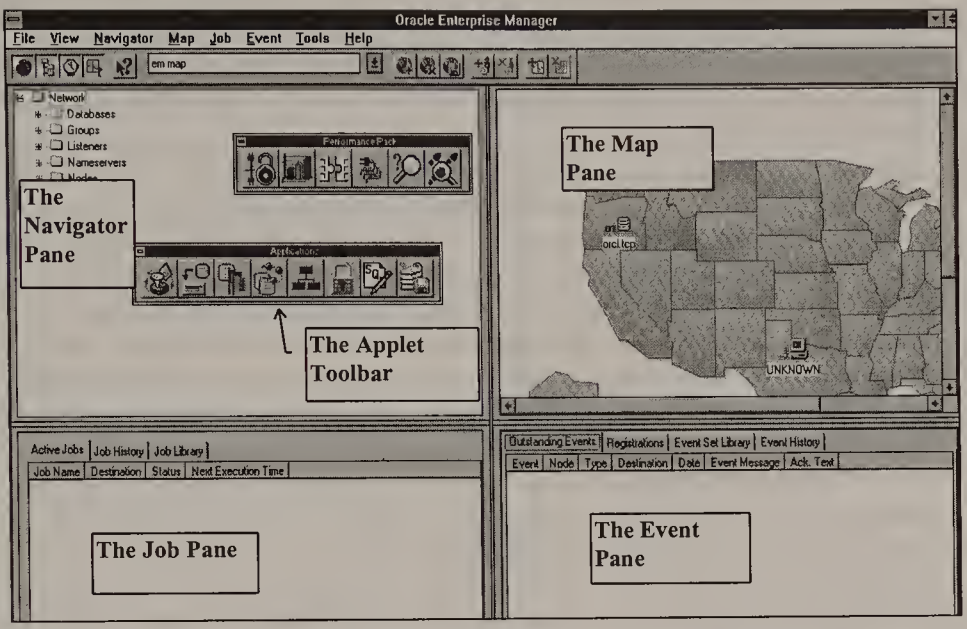


FIG. 13.4. Oracle Enterprise Manager windows console.

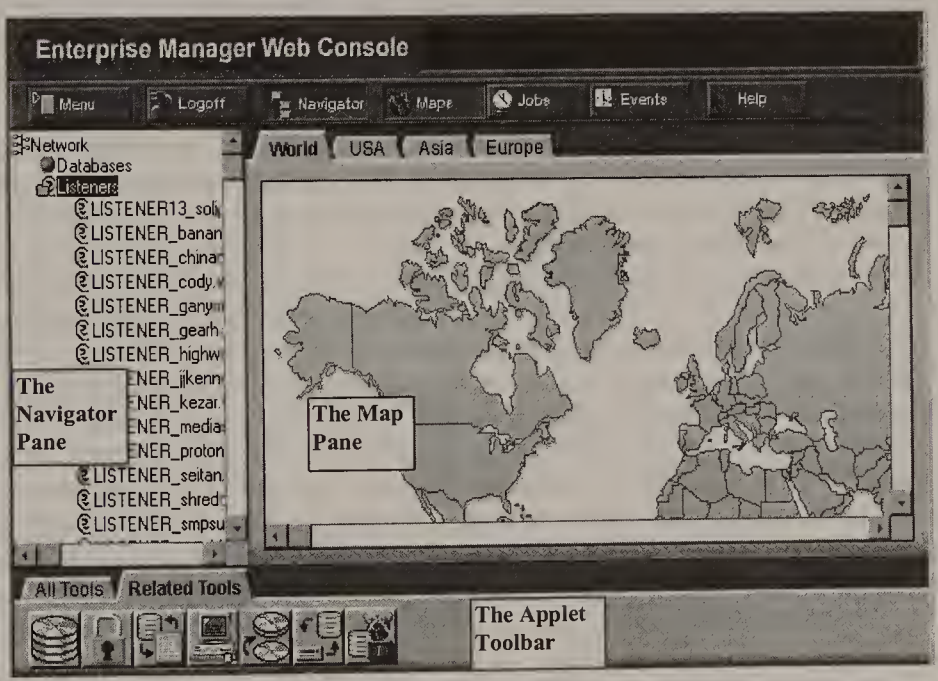


FIG. 13.5. Oracle Enterprise Manager web console, displaying only two panes.

Simplicity of a Single Window. Web applications running in a single browser window are more limited in functionality, but this seems to improve user performance over comparable non-Web multi-windowed implementations. The single window constraint is breaking down with frames. Thus far, developers have provided the basic functions that users want in a single window, without a lot of minor features they seldom use anyway, and users' performance appears better and learning faster than with multiple windowed applications. The HotJava Views™ user environment for NCs developed by JavaSoft, Sun Microsystems, Inc., also uses a single window approach (Gentner, Ludolph, & Ryan, 1996).

Negative Transfer of Training. Certain interaction paradigms (e.g. double click activation, right mouse click to pop up a menu, drag and drop) die hard among PC users making a transition to the Web. In fact, in our NC study, non-Web business users highly familiar with Windows and Macintosh software required the most assistance from testers of any user group. With Java programming, many of these interaction paradigms can be supported.

Web Applications Starting to Look Like Windows. Many Web software developers are gravitating toward a traditional PC appearance in Web applications. Looking like Windows but not having a comparable response time or interaction paradigm may further aggravate negative transfer of training and disappoint user expectations. If a new UI paradigm is to emerge independent of current platform conventions, style guides showing novel UI designs are needed.

Web is Slow. Web response time is a big problem even on LAN-based systems. In a survey of 23,300 Web users, 69.1% said the biggest problem was that it took too long to view or download pages (Pitkow & Kehoe, 1996). Non-Web users in our studies immediately reacted to the slowness of the Web the first time they tried it. A recent UI article discussed a target time of 15 seconds for page loading (Nielsen, 1995). This time exceeds by a factor of 7 Shneiderman's (1980) 2 second recommendation for simple response times, and may still be unachievable by most Web pages, because of future increases in Internet traffic and Web page design complexity.

Web Content Not Motivating. The content of the WWW leaves many end-users cold; they see it as boring, not entertaining, with topics of little interest to demographic groups other than white teenage boys (Kraut, Scherlis, Mukhopadhyay, Manning, & Kiesler, 1996). This is a UI concern. If the content does not motivate users to keep going, they are less likely to get past the learning curve and more likely to drop Web use altogether.

Problems With Page Metaphor. Web novices do not understand the page metaphor; they do not assume they can scroll down, because they do not know they are on a page. They also do not understand the time it takes to load pages from a server.

Unexpected Data Loss. Users do not expect the consequences of navigating off the page where they were entering data. They typically get distracted by advertising or a promising link while they are filling out a form and navigate away. When they come back their input is gone. A more advanced

caching mechanism may be necessary to prevent this in future devices, or at least a warning that users will lose what they enter if they navigate away.

Confusion With Browser Affordances. Users typically do not differentiate between a content page or application and the browser. Applications should not rely on browser buttons to complete interactions. Not all browsers provide the same affordances.

Television Implementations Need Different UIs. Most current Web pages are unreadable on NTSC displays and cannot be successfully used with a remote control, especially if extensive typing is required. Colors and fonts need to be designed specifically for the limited range of television, and most interaction models used on PCs need to be simplified. A recent survey showed that 52% of home users would like to work with the Web on a television (Seminerio, 1996). Designers of web pages and applications should consider television input device and display limitations during development, instead of trying to convert from a PC UI after the fact.

FUTURE WORK

Our work on enabling greater levels of user interaction with the WWW has really just begun. We think that we have made only a few tentative steps toward exploring how users will interact with this vast new medium of data communication. Oracle is making a major investment in developing technologies and applications that will enhance users' abilities to do useful work and engage in commerce on the WWW. In particular, the following efforts are underway:

Tool products development. Database tools such as Oracle Discoverer™ and Oracle Forms™ are being adapted to run on the Web. These tools will allow users to explore and query Oracle databases and interact with database forms from any Internet browser, with full functionality equal to Windows 95 products. We are studying user requirements for these products and carefully noting differences that emerge from standard PC client applications' interfaces to prevent user problems with the Web implementations.

Internet classroom testing. We are developing a new technique in usability testing in our Internet classroom which bears similarities to the audience evaluation methodologies used by major entertainment studios. Multiple subjects are run simultaneously. Data is collected using on-line forms on the Web.

NC applications UI development. Productivity applications such as word processor, spreadsheet, and presentation graphics that run on a Web server are under development for future versions of the NC. Other products at Oracle, such as InterOffice, are also becoming NC-friendly.

REFERENCES

Ellison, L. (1996, January 19). Oracle Open World keynote address, *Oracle Open World.*

Gentner, D., Ludolph, F., & Ryan, C. (1996). Designing the HotJava Views™ user environment for a network computer. http://www.javasoft.com/products/hotjava views/hjv.white.html (accessed October 29, 1996).

Hof, R. D. (1996, June 24). Why corporate America is leading the charge. *Business Week*, 78.

Kraut, R., Scherlis, W., Mukhopadhyay, T., Manning, J., & Kiesler, S. (1996). HomeNet: A field trial of residential Internet services. *Communications of the ACM*, *39*(12), 55–63.

Nielsen, J. (1995, December). Guidelines for multimedia on the Web. http://www.sun.com/current/columns/alertbox.

Oracle Corporation. (1996, September). Network Computing Architecture: An Oracle white paper. http://www.oracle.com/nca/html/nca_wp.htm

Pitkow, J. E., & Kehoe, C. M. (1996). Emerging trends in the WWW user population. *Communications of the ACM*, *39*(6), 106–108.

Seminerio, M. (1996, October 31). Survey finds users want their Internet TV. http://www.news.com.

Shneiderman, B. (1980). *Software psychology*. Boston, MA: Little, Brown & Co.

Wichansky, A. M., Friedland, L., Slater, P., & Dirik, A. (1996). Interdisciplinary software design. In *Proceedings of the Silicon Valley Ergonomics Conference and Exposition, ErgoCon'96* (pp. 347–350), Palo Alto, CA: San Jose State University and Silicon Valley Ergonomics Institute.

14 Web Usability Research at Microsoft Corporation

Amy Kanerva
Kevin Keeker
Kirsten Risden
Eric Schuh
Mary Czerwinski
Microsoft Corporation

The Interactive Media Usability Group at Microsoft Corporation covers a broad range of Microsoft divisions and product categories. Our group covers Web content, reference products, games, desktop finance products, children's products, and input devices—well over 100 different products in all. Although we have a fairly large group of researchers (10 full-time and 11 part-time employees), covering so many different product areas challenges even the most knowledgeable usability engineer. Sharing empirically based principles and design guidelines across team members is critical as development cycles shorten.

As usability engineers at Microsoft, we support the design process with traditional lab and field research techniques, as well as with other, more "streamlined" usability activities (such as participatory design, heuristic evaluations of preliminary specifications, and testing of low-fidelity paper prototypes). Specific usability goals, identified at the beginning of the product cycle, must be met before the design is frozen. To ensure meeting these aims, usability goals, milestones, and research activities are built directly into the development schedule.

This chapter reviews the development of guidelines and principles for three Web research areas that we have found to be critical: organization of the content, navigation, and engagement. Each section of the chapter touches on important questions and varied research methodologies used to address one of these research areas.

Early in the design cycle, we perform research to understand the optimal organization of the Web site. The first section of this chapter describes how we collect card-sort data from users and analyze the data into an organized, conceptual Web site model by using hierarchical cluster analysis.

After the optimal organization has been designed from a user's perspective, our usability engineers working on Web products turn their attention to navigation issues. Identifying the structural design properties that support effortless navigation through the site can be worked on very early in the design cycle, often with flow charts or paper and pencil prototypes. Because of the vast

number of ways to research this aspect of a Web site, the second section focuses most heavily on the principles and guidelines that have emerged from this work.

Simply having a good organization and structured design, however, is not enough to ensure a successful Web site. Once the foundational design of a Web site is intact, a usability engineer needs to ensure that the content is appealing. The chapter closes with a discussion of our research efforts for an engaging experience on the World Wide Web (WWW).

Each of these three areas of Web design is discussed in the order in which it occurs in the design process at Microsoft. For each section, we discuss the motivation behind the usability work, the general methodology developed for this topic area, and a selected set of Web design guidelines. Even though this chapter does *not* focus on *all* of the standard methods used by our group throughout the development of Web products (e.g., paper prototyping, heuristic evaluation, iterative lab testing), these methods are used heavily by usability researchers at Microsoft and have been critical to the guidelines effort.

ORGANIZATION

Motivation

Taxonomies are ubiquitous in everyday life—the division of movies into drama, comedy, and horror at the video store; the intricate Dewey decimal and Library of Congress classifications at the library; the placement of foods in the grocery store. A well-constructed taxonomy is a boon to end users, and enables them to quickly and accurately locate desired information. Conversely, a poorly built taxonomy can lead to a considerable expenditure of time and effort that fails to produce the desired information.

There are several approaches to building a taxonomy. The *organic* approach can rely on a combination of naive intuition and incorporation of existing structures. On the face of it, the Microsoft Network (MSN) version 1.3 (http://www.msn.com) and America Online™ (AOL) version 3.0 (http://www.aol.com) online services both seem to be examples of this organic approach. The *expert* approach relies on professional indexers to build and code a taxonomy. Yahoo's (http://www.yahoo.com) categories are the clearest example of this expert approach. The *user-centered* approach incorporates user feedback in both the construction and validation of the hierarchy. Because the taxonomy is built—and tested on—end users, the user-centered approach has a higher probability of mapping with common mental representations than does either of the other approaches used in isolation. The result is a hierarchy that is easier to use.

The first step in building a user-centered hierarchy is developing an understanding of how people organize information space. This understanding can be gained by using hierarchical cluster analysis to analyze data collected via card-sorting tasks. The resulting tree structures can form the foundation for organizing the site. Moreover, it is possible to identify meaningful patterns in the resulting

hierarchy which are indicative of general underlying cognitive processes. These patterns can be abstracted to general principles and guidelines for organizing Web content.

General Approach

Early in the development cycle, long before any content has been created or any interface has been designed, it is possible to do research that dramatically impacts the overall usability of the site. Card-sorting studies can be conducted by using little more than the names of the site elements. From these humble origins, a strong, user-centered organizational structure can be built.

Typical card-sorting studies involve 40 to 50 participants. Researchers may use different populations that suit specific research questions. For example, children can help to investigate the organization of kids' Web sites, and the hierarchies generated from novices and experts can be compared for underlying similarities and differences to aid design.

Participants in these studies are typically given a number of index cards, each printed with a description of a particular part of a Web site. When testing the organization of existing Web pages, it is possible to use a printout of the pages instead of a written description. Participants are asked to sort the index cards into meaningful or related piles and are often asked to provide a label for each pile. Comparing labels to the clusters that emerge from analysis can facilitate interpretation.

The results of the sorting task are then entered on a spreadsheet. Items that occur together in a group are given an identical category number. These sort values are used to build a frequency of co-occurrence matrix across all the subjects on which various cluster analysis techniques can be applied. Each entry (i.e., each item described on a card) starts out as a single unit or cluster. Additional entries or clusters are iteratively joined to it on the basis of a similarity criterion. The end result of this process is a tree-diagram that provides a view of the number of clusters and the content of each cluster.

Another variant of this approach can be used to test the ease of use of a proposed organizational scheme. In this version, participants are given a printout of the suggested hierarchy, as well as cards containing descriptions of the content. Participants are then asked to replace the cards into the hierarchy as they see fit. Participants may also be asked to rate their certainty about placing the card in the "right" place. Various "goodness of fit" metrics can be calculated (e.g., percentage of people who categorize the item the same way the originator of the hierarchy did, average certainty ratings), and these can give a sense of how good the proposed hierarchy is. It is also possible to develop a better understanding of how people represent the information space, especially when there are common "misclassifications."

Guidelines

Several cluster studies have been conducted to help organize information available on MSN, for version 2.0 of the product. These studies focused on the

organization of Web content, proprietary content, and children's content. Although many results are too specific to be of broad interest, several recurring principles emerged:

1. *People pay attention to the source of information, at least as much as they pay attention to the content itself.* As a result, the University of Washington Psychology page (http://weber.u.washington.edu/~psychol) tends to be classified with other academic pages (even those from completely different departments) rather than with other psychology pages.When tagging Web content for search mechanisms, designers should make sure to tag the source information. Designers of search mechanisms might want to group search results by source.
2. *People make the most classification "errors" at the deep levels of the hierarchy, when they are asked to make fine distinctions they may not completely understand.* Designers should reduce the number of low-level categories and make sure they are readily discriminable.

3. *People have difficulties classifying objects that exist at different levels of abstraction.* For example, people classify the New Yorkers for Ralph Nader Web page (no longer available) as applying to politics, U.S. politics, presidential campaigns, local politics, politicians, and celebrities. Items that exist at multiple levels of abstraction should be classified at each level of abstraction.

4. *A certain percentage of the population simply does not have the requisite world- knowledge to locate a site.* For example, 15% of people in one study believed that Madagascar was in Europe. Designers should consider incorporating common misconceptions into a hierarchy even if they are not "correct."

After the basic organization has been developed by taking into account user feedback, usability attention turns to navigation.

NAVIGATION

Motivation

Easy navigation is a critical component of a Web site's success. Even if the basic organization is clear, a poorly designed navigation structure causes users difficulty. Inconsistent navigation may also affect perceptions about a site's content. Several studies have shown that engagement is positively correlated with ease of navigation. Thus, building clear navigation is a key factor to immersion and engagement.

Designers of Web navigation schemes often fail to provide distinctive landmarks that indicate where users are located and where they can go. Real-world navigation often involves following clear paths or trails, but Web sites frequently do not provide maps or trails to guide movement through Web environments. At an extreme, the only way users can tell where they are in some Web pages is to remember where they came from. This process can be like trying to remember, in order, a long list of random numbers.

To help users to develop a sense of a Web site's space, the navigation design needs to systematically add cues and provide structural elements that serve as spatial landmarks and navigational devices. Landmarks and navigational devices are particularly important in designing Web sites for children as they are less likely to have general knowledge about spatial navigation. At the same time, when strong landmarks are absent, adults tend to revert to less developmentally advanced ways of reasoning about a space. Our goal in a series of recent studies was to understand how to engineer fluid navigation into Web sites.

General Approach

Usability activities focusing on navigation tend to occur relatively early in the design of our Web products, although competitive analysis can support this process at any stage of design. One goal of our research was to extend our existing navigation design guidelines that we derived from spatial cognition and design literature. The extensions to these guidelines would be based on our own studies with Web users, by conducting usability research that helped us establish a basic understanding of which cues seem to be most relevant for Web site design. The scope of this research relied on successive usability studies with a combination of low-fidelity paper prototypes, higher fidelity online prototypes, and existing online service providers such as AOL™ (http://www.aol.com) and Pathfinder™ (http://www.pathfinder.com). Across multiple studies, data about behaviors during use and data about postuse memory of the Web sites were collected. Different studies included various dependent measures (e.g., user errors categorized according to a specific coding system [see Table 14.1], search times, and users' memory of their traversal paths, analyzed for accuracy and completeness).

Guidelines

Web site design should facilitate navigation by using visually salient landmarks in

TABLE 14.1
Error Classification Scheme Used During Web Usability Studies

Type of Error	*Behavioral Indice Observed*
Where am I?	Behavior indicates users do not know where they are in the structure of the Web site.
How can I get to other places?	Behavior indicates users do not know how they can navigate to another place from the Web site.
Where can I go?	Behavior indicates users do not know where they can go in the Web site.

the form of labels, buttons, headings, and icons. These visual cues should indicate the current location and potential pathways to clearly described ultimate goals. The following guidelines provide an overview of design techniques and elements that can be used to support spatial learning and navigation. All these guidelines have been empirically validated in the laboratory with target end users.

1. *Simplify the virtual "landscape" for users so that spatial information can be easily learned.* Designers should not overload pages with choices; as a general rule, five choices per page are sufficient, unless they can be grouped as a single chunk of information. Designers should make all choices in an image map fit in one screen. Directories and table of contents pages must have intuitive and recognizable ways of grouping and organizing entries. If a navigation bar contains both links to places and access to tools, the design must highlight the difference between these button groups. Designers should avoid nested scroll bars, and should especially avoid placing a nested scroll bar next to a major scroll bar. When using scrolling, designers should signal that there is additional content below or to the side.
2. *Make users aware of where they are at all times.* Metaphors can be helpful in communicating the organization of a space. When using a metaphor, carry it throughout the navigation design. Otherwise, the lack of consistency actually impedes its effective use. Navigation bars should be very noticeable and should communicate which category users are in. Every time a site is mentioned, it should be given the same name, and each page should be labeled with its topic. Users should recognize changes as they navigate from one page to the next.
3. *Make users aware of where and how to travel in the space.* Create a persistent navigation structure to guide users to a home base at the major, intermediate, and minor levels of the Web site organization. Never assume that users can use the Back button to get out of pages; at the same time, always assume that some of them will use the Back button instead of the button you have created to get them out, and plan for this outcome. Never assume that all readers arrive at a page from the same place. Avoid navigational dead ends and links to incomplete pages. Place icons and buttons in a location that takes into consideration where attention is likely to be drawn.

   Make sure to distinguish graphic content images from navigational icons and buttons. Avoid intrapage links. If intrapage links must be used, ensure that users can see a link back to the top of the page. Make sure that navigation buttons and icons are meaningful for target users. Make sure icons that are clickable in one location are clickable throughout the site, unless they are grayed out. Make sure that a particular icon always takes users to the same place wherever it is encountered. Avoid ambiguous terms, if possible. Use labels that indicate specifically where a navigation element takes a user. Use tool tips to reinforce and clarify functions and destinations of buttons and icons. Build preferred paths into designs (rather than using a tour or guide).

Preferred paths may be important for engagement with some Web sites (e.g., narratives, productivity).

Once designers know the basic organization of the Web site and have successfully designed in the ways users will navigate through it, they look at the content. We have been creating content geared toward entertaining home users. The following section focuses on how we approached this research.

RESEARCHING APPEALING WEB CONTENT

Motivation

At the end of 1996, the Microsoft Network discontinued its use of proprietary protocols, moved to the Web, and delivered all its content via standardized Web protocols, common scripting languages, and ActiveX controls. MSN version 2.0 also kicked off a movement to broaden the appeal of online services by bringing more entertaining content to the Web. Like television, theater, and movie entertainment, these new content offerings were called *shows*.

Web shows create a special challenge for usability researchers. Although shows must be usable, typical usability goals like simplicity and consistency may not always be appropriate in an entertainment context. People want their games to be challenging, their books to be frustratingly suspenseful, and their tearjerkers to make them cry. Usability researchers must try to keep interface problems from lessening the impact of the content while realizing that content and user interface (UI) are inextricably tied.

Shows also present a measurement problem for researchers. The ultimate success of entertainment software is difficult to measure in tangible, reliable behaviors like task time or number of errors. In addition to traditional measures, researchers have to make strong use of natural observation and subjective questionnaires. The ultimate measure of success for a show is whether people enjoy it, become engaged in the experience, and return to the show again and again.

General Approach

One part of Microsoft's approach to show research has involved creating a reliable, subjective questionnaire to measure appeal. In doing so, we have tried to identify a set of factors that contribute to the appeal of Web content.

Our observations have been drawn from many types of entertainment venues. Users look for particular characteristics in certain types of content. For example, electronic magazines need to be timely, while networked Web games need to accurately target a range of skill levels. We assume, however, that the lessons learned in one domain may be, in some way, applicable to another domain. For example, we expect soap operas to have strong, identifiable characters, but we do not expect reference Web sites to have strong, identifiable characters. Interestingly, movie reviewers play both a reference and an

entertainment role. A reference Web site with a strong editorial viewpoint, in addition to an unbiased perspective, could be very appealing.

Factors relating to appeal were iteratively derived via a combination of theoretical and empirical means. Theoretical insights came from researching the Human–Computer Interface (HCI) and psychological literatures and from frequent discussions about entertaining media. Empirical research focused on kids' entertainment titles and age nonspecific entertainment content for the WWW.

In an ongoing Microsoft research project, children rated their favorite Web sites or computer products on Likert-style questionnaires (younger children marked their ratings on a graphic response bar with iconic positive and negative poles). Usability engineers derived items from Lepper's (1988) research on enhancing intrinsic motivation, Whalen's and Cziksentmihalyi's (1991) research on flow states, and UI research by Hix and Hartson (1993). A factor analysis of the questionnaire verified Lepper's (1988) suggestion that challenge, control, contextualization, and curiosity are important underlying dimensions of intrinsic motivation.

A literature review of HCI research on games suggested additional items to assess: ease of use and attractive use of graphics and sound (Carroll & Thomas, 1988; Igbaria, Schiffman, & Wieckowski, 1994; Malone, 1981; Neal, 1990; Pausch, Gold, Skelly, & Thiel, 1994). Concurrent discussions about entertaining media focused on emotional reactions to stories told via television, movies, and books. A range of primary emotional reactions was pulled from the psychological literature on emotion and added to the questionnaire. User comments and market research suggested that users returned to those Web sites that had the correct amount of quality information on a topic interesting to a particular user.

At each stage, item sets were analyzed empirically. Weak questionnaire items were removed or revised to leave a reliable core item set. The resulting item sets were used to research MSN shows. Systematic analysis of data collected during iterative usability testing, site visits, and broad distribution surveys over a year has recently shown that most of the attributes listed in the following Guidelines section predict how much people like a show and how many hits a show gets once it is released. To further validate the appeal factors, we expect to compare appeal questionnaire ratings to measures of actual MSN show usage rates.

Guidelines

A structured series of guidelines for enhancing appeal was derived over the course of extended research and refinement of the questionnaire. Each guideline may not be applicable to every show. Some guidelines are more important for certain types of entertaining experience than for others.

1. *Ease of use—clear goals.* The goal of the show is immediately clear, and the steps to achieve the goal are clear.
2. *Ease of use—clear structure.* The design clearly communicates the show's core activities and sets appropriate expectations. The terms (especially the title

and the sections of the show) adequately communicate the consequences of selection or action. The core activities require few actions to locate. People always have more than one course of action to choose from. People are able to get out of sequences that they do not want to be in.

3. *Ease of use—clear feedback.* Page headings help users to understand where they are located. Design elements (e.g., headings, scores, visual metaphors) indicate what is happening.
4. *Content quality—valuable.* The topic matter is interesting to the core audience, and people have an opportunity to learn.
5. *Content quality—timely information.*The show features the latest information available on the topic.
6. *Content quality—appropriate depth and breadth.* The show has the right amount of variety for the topic area. The articles are the right lengths, and the show goes into the appropriate level of depth in its topic areas.
7. *Content quality—attractive use of media.* The graphics and audio are appealing to the core audience.
8. *Content quality—safe novelty.* Users experience a mix of the known and unknown, in a novel twist on the familiar.The metaphors, language, graphics, and sounds set appropriate expectations for the show.
9. *Content quality—challenging.* The level of challenge in the show's content is appropriate for the core audience. People find new challenges as they gain experience.
10. *Content quality—evokes emotion.* Users feel emotion while experiencing the show.

SUMMARY

Designing the organization, navigational structure, and appeal of a Web site from the perspective of users can be a difficult process because of the large number of Web sites and the rapid development cycles with which usability engineers typically work. It has been our observation that pulling guidelines and principles from the empirical data we collect during Web interaction can have positive impact on Web site design. This chapter has presented three of the many important areas of Web design and has shown how the Interactive Media Usability Group at Microsoft is involved during these design stages. In the end, the usability group needs to work to ensure that users come back to our Web sites again and again. Our research efforts continue to focus on Web designs that evoke emotion, interest, and pleasure for our target audiences.

REFERENCES

Carroll, J. M., & Thomas, J. C. (1988). Fun. *SIGCHI Bulletin, 19*(11), 21–24.
Hix, D., & Hartson, R. (1993). *Developing user interfaces*. New York: Wiley.

Igbaria, M., Schiffman, S. J., & Wieckowski, T. J. (1994). The respective roles of perceived usefulness and perceived fun in the acceptance of microcomputer technology. *Behaviour and Information Technology*, *13*(6), 349–361.

Lepper, M. R. (1988). Motivational considerations in the study of instruction. *Cognition and Instruction*, *5*, 289–309.

Malone, T. W. (1981). Heuristics for designing enjoyable user interfaces: Lessons from computer games. In *Proceedings of the Conference on Human Factors in Computer Systems* (pp. 63–68). Gaithersburg, MD: ACM.

Neal, L. (1990). Implications of computer games for system design. In D. Diaper (Ed.), *Human-computer interaction—Interact '90* (pp. 93–99). New York: Elsevier.

Pausch, R., Gold, R., Skelly, T., & Thiel, D. (1994). What HCI designers can learn from video game designers. In *Proceedings of computer human interaction* (pp. 177–178). Boston: ACM.

Whalen, S., & Csikszentmihalyi, M. (1991). *Putting flow theory into educational practice.* (Report to The Benton Center for Curriculum and Instruction). Chicago: University of Chicago Press.

15 Creating Content for Both Paper and the Web

Richard C. Omanson
Gavin S. Lew
Robert M. Schumacher
Ameritech

The World Wide Web (Web) has become a legitimate medium that businesses use for advertising, sales, and customer support. For most companies, the question of attracting and keeping new customers is no longer whether to get on the Web, but rather what type of presence to have. Similarly, many companies are finding it cost effective and useful to make internal corporate data available electronically on an internal internet (or intranet). Starting from a green field—without having to deal with old information—would offer no problem; all data would be tuned primarily to electronic presentation and secondarily to paper artifacts. But beginning from scratch is a luxury that few companies have or can afford; businesses find themselves with considerable investments in legacy data and media. Until recently, the electronic world was tuned to produce print artifacts. Perhaps more important, people's thinking in production and distribution of documents was geared toward producing print. What we argue here is that to be successful in publishing in two mediums, the zeitgeist must change.

Technological and competitive forces make people look squarely at the problem of how to approach integrating Internet delivery of information with existing paper documents and print-based collateral. In most cases, organizations do not simply switch wholesale from paper to the Web: The costs are too high, the resources are limited, and the tools for doing this conversion can be daunting. Rather, companies maintain and develop both types of media. Our experience is that organizations differ in what they think is involved in using both Web and paper to perform a function such as advertising. It is just these differences that can be characterized in three distinct stages: paper conversion; joint development; and Web optimization. Looking at an organization's effort in terms of these stages helps us understand what issues they currently have and what issues they are likely to have in the future. This chapter will describe these three stages of maintaining Web and paper media, the issues that arise in each and research we have done to address some of the issues. To date, our research has concerned issues arising in the first stage.

PAPER CONVERSION

Description

In the paper conversion stage, creating content for the Web is perceived to be the "simple task" of converting from existing paper to electronic media. In this stage, getting content on the Web is a two-step process: First, the material is developed for (or acquired from) a paper medium and then the same content is recreated for the Web. The reason for the primacy of paper in this stage is easy to understand. Managers placed in charge of producing a Web channel typically have already developed effective channels in paper media such as magazine ads or employee newsletters. When they are given the task of creating content for a Web site, they do not have the budget or the time to start from scratch. Inexperience with Web pages also leads to the assumption that the capabilities of paper and Web are largely similar. Thus, managers take existing paper documents and simply "put them on the Web."

Paper Conversion Issues

When an employee or vendor is given the job of putting an existing content in electronic form, many unanticipated difficulties often appear. These difficulties arise because as a medium, paper has different constraints from the Web.

The first area of concern is the text itself. The conversion process often involves simply taking text from paper and dumping it into an electronic form, with the assumption that the readability of the paper and electronic documents is the same. Studies have found no difference in reading comprehension between paper and electronic documents (Dillion, 1994), but, content written for paper does not necessarily read as well in electronic form (see Gould & Grischkowsky, 1984; Wright & Lickorish, 1983 for a thorough review). When converting paper documents, it is prudent to assess the electronic document's readability. If there is concern that the electronic document may be difficult to read, then the electronic text should be changed.

One way to improve the readability of electronic text is to change the content so that readers essentially have to read less. This task is nontrivial because careful attention must be given to the conversion process: Reducing the amount of text can also reduce important information. Several electronic document authors suggest writing style techniques that can improve readability of electronic documents (Horton, 1990, 1991; Wieringa, Moore, & Barnes, 1993).

For example, at Ameritech and most companies, standard methods and procedures documents (M&P) are written in paper. These M&P documents include service or product descriptions and technical information. These documents are very large (so large that documents are often measured in inches instead of pages). When one business unit decided to convert its M&P documents to an electronic medium, the readability of the converted text was a concern. We brought in experienced electronic writers and technical editors to assist us in our conversion process. Writing styles best suited for electronic documents were taught and implemented. To the pleasure of both authors and readers, conversion

to an online reference system produced a more concise and useful M&P document through the standardization of content, changes in writing style, and merging of unique references previously unavailable to M&P documents (e.g., sales and marketing materials, such as bill inserts and advertising promotions).

The second area of concern when converting paper documents to the Web is that often the graphics and photographs used on paper cannot be used on the Web. In the paper medium, the file size of the graphics or text on the page has no impact on user experience. On the Web, the text and graphics are downloaded from the visited Web site and then assembled as a coherent page. Here, the file size of the graphics and text directly impacts the user experience. When users have to wait 2 to 3 minutes for graphics to download, the Web medium imparts more irritation than information. If the graphics on a site are too large, users set their browsers so that the graphics on the page are not included. Because the file size of graphics has no impact on the experience of readers in the paper medium, it is often the case that the file sizes of the graphics and photographs used in paper documents are simply too large to be used on the Web.

For example, in 1994, we converted glossy collateral about distance learning products as part of an education solutions section of our Web site (this content has since been retired). Two thirds of the photos in the collateral had to be eliminated, both because their image size was larger than most screens and because their file size was too large (approximately 80K) for reasonable download times. Several diagrams were used, but they were not very satisfying. The diagrams took up the entire screen and could not be reduced without making the captions and explanatory text illegible.

As a third area of concern, page layouts used in the paper medium require more resolution than is available on a computer screen. For example, a legible font size on paper media may allow for three columns of text, but these same three columns of text do not fit on most Web pages unless special graphics are used. A related effect of less resolution is that text is harder to read on computer screens than on paper. Some studies have found that reading speed of paper documents is 20% to 30% faster than online documents under common conditions (e.g., Gould & Grischkowsky, 1984; Muter, Latremouille, Treurniet, & Beam, 1982; Wright & Lickorish, 1983). Although improvements in the quality and resolution of computer screens over the past decade have reduced this disparity, resolution advantages still remain for paper documents (Crown, 1991). As a result, pages with high text density do not work as well on the Web as they do in print. For example, in Ameritech's 1994 Annual Report, in a section describing our strategy of being the best full-service communications company, descriptions and photographs of five key products are presented on a single page in three columns. When this section was converted to the Web, much of its punch was lost because these product descriptions had to be put on separate pages (www.ameritech.com/news/annual/html/strategy1.html).

A fourth issue in converting content from a paper medium is navigation. With paper documents, an author can specify the size and layout of each page of a document and the way in which the information is presented in and across pages.

And because the amount of information on a page is controlled, the navigation through a document is always the same, to wit, page turning. On the Web, authors do not know how much of a "page" is visible to a given user because screen size, window size, and screen resolution are all determined by a user's computer and browser settings. Because the page is not under the control of the author, but is controlled by a user's screen and browser settings, content organization (links to related topics) and explicit navigation bars, rather than page organization, must be used as the basis for navigation. Aside from page layout, navigation in print is made available primarily through tables of contents and indexes. If anything, a table of contents in electronic form is more usable because it can be made into hypertext. The notion of an index is an odd one electronically. Most index-like functions are made available through the browser on a Find or Search feature, which reduces to nil the possibility to serendipitously find information through an index.

A fifth issue in converting content is context. In a paper document, there is minimal skipping around so that authors assume that previous pages have established a context for subsequent pages. On the Web, search engines, bookmarks, and advertised universal resource locators (URLs) enable users to go directly to a page without first going through previous pages. The impact of this direct navigation is that users bypass the context designed into the site and become "lost in hyperspace" (Dillion, 1994; McAleese, 1989). To reduce this difficulty, each Web page must provide enough context that users can quickly determine where they are in the document.

Paper Conversion Research

To address some issues of context, we began a program of research about the cues contained within documents to establish context. In print media, headings are used to reveal document structure. For example, a book contains chapters, followed by sections in each chapter. Several researchers have found that headings provide structure aiding reading, recall, search, and comprehension in paper documents (e.g., Hartley & Trueman, 1983; Spyridakis & Standal, 1987; Wilhite, 1991).

Most authors and publishers have related the importance of the content through such structural elements, but intimately correlated with structural headings are the appearance elements such as boldface, underline, and centering. For example, an author or editor might use differing font size, boldface type, underline, or italics to help readers distinguish between a chapter on animals and a subsection on mammals. Because it is possible, and often happens, that a reader winds up in the middle of a text that the author intended to be read serially, Lew and Schumacher (in preparation) questioned whether certain characteristics about a heading might reliably signal the relative level. If there are such characteristics, using them might help people recognize these headings and orient themselves when landing in the middle of a document from a search engine or bookmark.

Two experiments were done to assess user preference for heading-level formatting. In Experiment 1, participants played the role of authors and selected format characteristics (e.g., font size, boldface, italics, underline, and center) for

TABLE 15.1
Selected Format Characteristics by Heading Level

Heading	*Level 1*	*Level 2*	*Level 3*	*Level 4*
12 pt	0%	0%	21%	58%
14 pt	4%	58%	75%	38%
18 pt	96%	42%	4%	4%
Bold	71%	58%	67%	29%
Italic	4%	4%	21%	58%
Underline	33%	17%	8%	13%
Center	46%	25%	17%	21%

headings in an electronic document. In Experiment 2, participants played the role of readers and made forced-choice heading-level judgments based on headers with predefined formatting. To the extent that there is unity in these measures, the results suggested a set of guidelines for formatting the headings of online documents.

The results are presented in Table 15.1. In both experiments, participants tended to select higher font sizes and centering with higher heading levels. Italics and underlining were used for lower level headings. Boldface was used on all but the lowest heading level. The most robust of these effects was for font size. Postexperimental interviews suggested that readability issues underlie these findings. Participants said that higher heading levels denoted greater importance, so they selected larger font sizes to "catch readers' attention." But, some caution is warranted for this heading-level–font size guideline because it could justify simply increasing the number of font sizes to match the number of heading levels. Horton (1991) suggested that users can distinguish only among four font sizes when reading an online document. Thus, there is a limit to the number of font sizes that can be used to identify heading levels.

When the actual combination of formats used was examined, only a few format combinations were selected for each heading level, a finding suggesting that there are some online heading-level stereotypes for certain combinations of formats. The most frequent format selection for header level 1 was large font size (18 points) and boldface. The most common format for header level 2 was large font size without additional formatting. The format preference for heading level 3 was medium font size (14 points) and boldface. For heading level 4, small font size (12 points) and italics were selected. Hence, on the basis of these results, four unique format characteristics can be associated with heading level.

Joint Development

The joint development stage evolves from the paper conversion stage. The thinking goes something like this: "Conversion from paper source is painful. Conversion tools often provide only crude approximations for what is desired. As I know I want to produce both paper and electronic copy, I'll factor them into my

production. It is too hard to produce the electronic documents once I have paper, so when I compose the documents for delivery, I'll be smarter about the formats I produce." The goal at this stage is to minimize the work of deploying content in two mediums by creating single source content that can be deployed in either the Web or paper medium. Conversion is not eliminated, but it is made as painless as possible by pushing it into a production-level problem.

A good example of this strategy is Ameritech's employee newspaper, *Vision*. When the layout of this publication was revised, its appearance in both print and Ameritech's intranet was planned simultaneously. Section icons were designed to look good in print and to look like navigation on the Web. For each issue, the number of columns was kept to two. When a sidebar was used, it was designed so that it could be a separate Web page without losing its punch. The file size of graphics was kept small and color palettes were chosen to look good in both mediums. The result of these considerations is that conversion to HyperText Markup Language (HTML) is painless and nearly automatic. All the content can be used in both mediums and has a clear mapping to comparable sections of the Web and print documents. Once the new issue is complete, it is sent to the printer and the Web converter simultaneously and usually appears on the intranet before it is in print.

For the most part, the issues encountered in the joint development stage are the same as in the preceding stage. The primary new issue is going from one medium to the other. Readers of the print version may want to see the Web version. Linking print to Web is easily accomplished by simply including the URL for the Web version in the copy of the print piece. More problematic are Web users who want a hardcopy of what they see. For one-page ads, relying on the print function of browsers is sufficient. Large documents with internal structure and subpages are more problematical. For these documents, all the HTML files can be assembled into one. Or, the assembled files can be converted to a printable format such as Portable Document Format (PDF) from Adobe Systems Incorporated (www.adobe.com) The problem, of course, is that creating separate print and viewing versions violates this stage's goal of minimizing the work involved in creating a Web version of content. Creating separate version, however, is also the first step toward the third stage.

Web Optimization

In the third stage, web optimization, the problem of creating content for both paper and electronic mediums becomes a problem of document creation. As one of the main attributes of the Web is to allow multiple styles of interaction, paper delivery is simply conceived of as another requirement. Managers no longer try to keep content presented in the two mediums the same. Rather, properties such as feedback, interaction, animation, and multimedia developed for the Web medium have no clear counterpart in print. In Stage 1, content was designed first for print and then converted to the Web; now the reverse is true.

When authors take advantage of the unique capabilities of the Web, the content becomes more engaging and interaction becomes much more important

Readers can participate with authors by feedback and mail forms. Interaction with the page is provided by Java from Sun Microsystems, Incorporated (www.sun.com) and other software modules. Multimodal experience is provided by video, sound and graphic animation, and the entire experience can be tailored to each visitor by collecting and using information from previous visits.

An example of the effects of adding Web-specific capabilities to a document can be seen in Ameritech's 1995 Annual Report. With this document, every effort was made to take advantage of the unique properties of the Web. Frames were used to present a persistent table of contents and glossary section that was always visible. Macromedia's (www.macromedia.com) Shockwave was used to animate graphs of financial results. The complete spreadsheets could be downloaded in Microsoft Corporation's (www.microsoft.com) Excel or tab-separated text formats. Finally, Progressive Network's (www.real.com) Real-Audio was used to provide an audio walk-through of the report's highlights. The effects of this attention were dramatic. In a review of annual reports on the Web in 1996, Cato (1996) noted in *Chief Executive:* "Ameritech, after taking first place for its print report, earns kudos for its electronic version, as well as a marked improvement from last year's disastrous effort" (p. 51).

Several issues arise when these features are used. If interaction with the author is employed, an author must be available to answer the questions. If multimedia is employed, the pages now must include training or instructions on how to use the features. Platform differences, which are largely taken care of by the HTML language with simple text documents, suddenly become an issue with audio and video files.

SUMMARY

Producing a document one time for electronic delivery is not a particularly difficult task. The real difficulties arise when the document needs to be stored, edited, and republished. Document management and sustaining production underlie the real problems.

The stages can be thought of in a simplistic document production model. Documents evolve from document creation (authoring and editing) to production (typesetting) and distribution. The stages roughly work backward from this model. The paper conversion stage involves tackling the problem at the source. That is, electronic copy is just a different way to distribute. The joint development stage is largely one of looking at the problem more from a production point of view. In other words, the document creation process does not change, but the way the document gets composed for delivery varies. Lastly, Web optimization pushes the thinking about electronic delivery into the document creation phase. The tables have been entirely turned. At first, electronic delivery was an artifact of paper delivery; in the end, because of the broad range of interaction possible, paper is simply one of many delivery media.

REFERENCES

Cato, S. (1996). Reports on the 'Net': More hype than achievement. *Chief Executive, 117*, 51.

Crown, J. (1991). *Effective computer user documentation*. New York: Van Nostrand Reinhold.

Dillion, A. (1994). *Designing usable electronic text*. London: Taylor & Francis.

Gould, J. D., & Grischkowsky, N. (1984). Doing the same work with hard copy and cathode-ray tube (CRT) computer terminals. *Human Factors, 26*, 323–337.

Hartley, J. & Trueman, M. (1983). The effects of headings in text on recall, search and retrieval. *British Journal of Educational Psychology, 53*, 205–214.

Horton, W. (1990). *Designing and writing on-line documentation: Help files to hypertext*. New York: Wiley.

Horton, W. (1991). *Illustrating computer documentation: The art of presenting information graphically on paper and online*. New York: Wiley.

Lew, G. S. & Schumacher, R. M. (in press). Towards a guideline for formatting characteristics of headings on online documents.

McAleese, R. (1989). Navigation and browsing in hypertext. In R. McAleese (Ed.), *Hypertext: Theory into practice* (pp. 4–61), Oxford, England: Intellect.

Muter, P., Latremouille, S. A., Treurniet, W. C., & Beam, P. (1982). Extended reading of continuous text on television screens. *Human Factors, 24*, 501–508.

Spyridakis, J. H. & Standal, T. C. (1987). Cognitive dimensions of reading comprehension. *Reading Research Quarterly, 22*, 285–298.

Wieringa, D., Moore, C., & Barnes, V. (1993). *Procedure writing: Principles and practices*. Columbus, OH: Battelle Press.

Wilhite, S. C. (1991). Headings as memory facilitators: The importance of prior knowledge. *Journal of Educational Psychology, 81*, 115–117.

Wright, P. & Lickorish, A. (1983). Proof-reading texts on screen and paper. *Behaviour and Information Technology, 2*, 227–235.

16 The Ten Golden Rules for Providing Video Over the Web or 0% of 2.4M (at 270k/sec, 340 sec remaining)

Chris Johnson
University of Glasgow

The WWW offers huge scope for the distribution of multimedia resources. Text, graphics, video and sound can easily be retrieved by users with only a minimal understanding of the underlying communications mechanisms. Unfortunately, the quality of web-based resources often dampens this initial enthusiasm. The production values and editing standards of video and audio material can be extremely poor. This is hardly surprising because designers have little or no advice about the delivery of such material over the web. Previous guidelines for word processors, databases and even computer supported cooperative work systems cannot easily inform the design of web based applications. The unpredictable nature of retrieval delays and high download latencies have created new types of usability problems. Existing experimental techniques cannot easily be used to investigate the dynamic nature of interaction with distributed resources. This paper, therefore, presents a number of guidelines that are intended to improve video retrieval over the web. These principles have emerged during the development of a number of existing sites. Without such guidelines, users will be condemned to tie up their machines for long periods while mediocre footage is retrieved across heavily loaded networks.

TIME AND THE WEB

The very success of the WWW has led to two different sorts of problems. Firstly, the rising numbers of users has created a demand crisis. Many servers simply cannot satisfy the number of requests that are being made for their information (Dix, 1987). Secondly, the increasingly rich, multimedia resources that are available over the world's computer networks have created a supply crisis. The introduction of sound and video material has massively increased the amount of information that must be transferred between remote sites. In consequence, retrieval delays have reached a stage where many sites suffer from a form of 'electronic gridlock' (Johnson, 1995).

It is important to emphasize that many of the standard results from human computer interaction cannot be applied to web-based systems (Gray, England, & McGowan, 1994). For example, previous studies have concentrated upon the effects of predictable delays or delays that have some fixed upper bound (Kuhmann 1989; Johnson, 1993; Teal & Rudnicky, 1992). In contrast, the delays that arise during interaction with the WWW are unpredictable; they do not have a fixed upper bound. In the absence of any substantial body of literature that can be applied to support web-based video material, this paper presents a number of principles or heuristics. These guidelines are backed both by experience in the development of web pages and by the results of limited empirical investigations. They are, however, partial. They do not cover all of the design decisions that must be made during the development of a video site. For instance, the principles do not address the technical direction of video material (see May & Barnard, 1995).They are, however, intended to support the development of web-based video pages in the absence of more detailed empirical evidence in this area.

A number of common concerns motivate the principles in this paper. Firstly, designers must consider the physical limitations of the underlying infrastructure that supports the presentation and distribution of video material over the web. It is little use providing material for a world-wide audience if most users have to tie up their machines for anything up to two hours while a large video file is downloaded from a remote site. Secondly, designers must consider the characteristics of their users' tasks. These have a profound impact upon individual attitudes to both retrieval delays and the production quality of video material. Finally, designers must consider the visual appearance of pages that present video resources. Designers must provide some indication of the production quality of a film if users are to identify valuable footage from the mass of poor quality material that is currently available on the web.

Principle 1: Avoid Video for Video's Sake

The first question to be faced by any designer that is interested in providing video resources over the web is: *why bother with video in the first place?* Too often, video clips simply contain "talking heads". There is little or no value added by moving away from a purely textual or graphical medium. In contrast, Fig. 16.1 illustrates the effective use of video by the Hunterian Museum's web site (see http://www.gla.ac.uk/Museum). These stills provide some idea of the protective value of a Roman shield. Unfortunately, they do not provide an accurate impression of the "three dimensional" protection afforded by the armor which is provided by the original video clip. Such information is difficult, if not impossible, to convey using text or even a static exhibit on the Museum floor.

Principle 2: Make Sure That Users Can Access the Video

Having decided that video is necessary to a web page, designers must determine *if users can actually access the material.* It is of little benefit digitizing large quantities of footage if users do not have access to appropriate viewers. Figure 16.2 illustrates the way in which helper applications can be called to play video

material by WWW browsers. Providing users with information about how to set up video viewers might seem like a trivial issue but, again, it is important to realize that this can have a considerable impact upon the effectiveness of video



FIG. 16.1. Excerpts from the Hunterian's video on Roman armor.

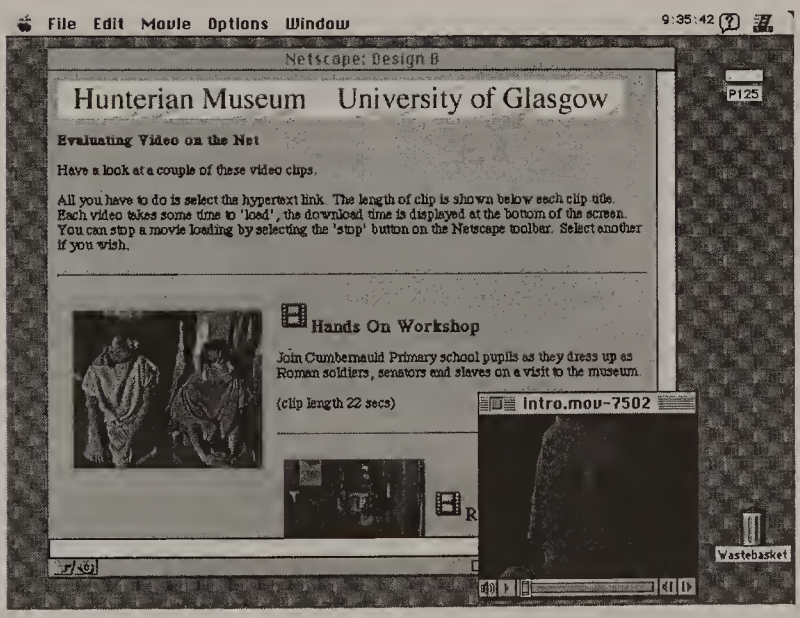


FIG. 16.2. Using a browser and a helper application to access video on the WWW.

material (Johnson & Kavanagh, 1996). At the very least, web sites that provide access to video resources must also provide links to appropriate viewers.A number of other factors ought to be taken into account if designers are to ensure that users can access video material over the web. In particular, the network infrastructure has a profound impact upon an individual's ability to download large files from remote servers. Figure 16.3 shows the amount of time that is required to retrieve the same video file under optimal loading using a number of different network technologies. The total running time for the clip was approximately eleven seconds. These findings suggest that designers must consider alternative media if the users of a video resource routinely require modem access. Under heavier loads, the delays become exponentially greater. As a result, the costs, both in terms of telecommunications charges and machine time, may be prohibitive.

Principle 3: Maximize the Production Quality

Users will only exploit video resources if the benefits of accessing the material outweigh the cost in terms of the retrieval delay. It is, therefore, important for web designers to maximize these benefits by ensuring high quality production. Techniques for directing and shooting video footage are addressed in a range of professional publications (Bernstein, 1988). However, production quality is not only determined by the way in which a film is shot. It is also affected by the techniques that are used during the digitization process. For example, frame sizes can be altered so that only a particular number of pixels are stored for each image. The lower the frame size, the lower the quality of reproduction. The number of frames per second also affects the quality of a film. The lower the number of frames, the lower the quality of movement. There are, however, no 'optimal' settings for these parameters. Films with rapid motion naturally require higher frame rates. Films that present close-ups of complex objects require higher frame sizes. The critical point here is that unless designers consider these issues then the resulting videos will be useless. It is no use down-loading a video of a horse running if the frame rate is so low that the user is presented with a jerky parody of the animal's motion. Similarly, it is of little benefit retrieving a video of a sculpture if the frame size is so low that the user cannot recognize the object. Gain and sample rates for sound reproduction play a similar role to that of frame size and rate for moving images. Unless designers consider these issues, then the poor quality of the video will have a direct impact upon the 'usability' of the final resource.

Principle 4: Edit the Clip to Maximize Content and Minimize File Size

After a video has been produced, designers must edit the raw footage into a number of clips that can easily be retrieved over the web. Unfortunately, this editing process cannot simply be based on artistic merit. Figure 16.4 (a) shows that download time is proportional to file size. Figure 16.4(b) shows that the

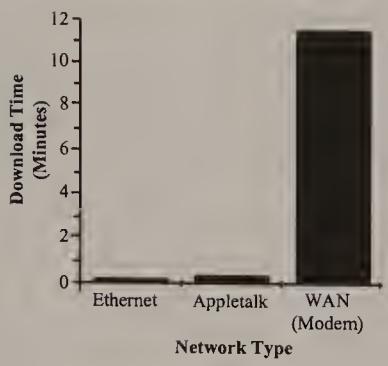
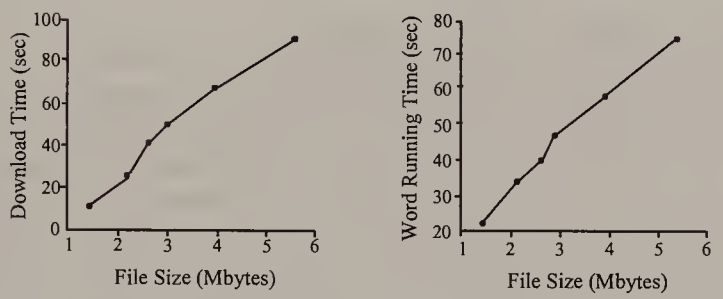


FIG. 16.3. Retrieval delays over different network connections.



(a) Variable Download Latency With File Size

(b) Variable File Size With Length of Video

FIG. 16.4. The relationship between video file size, download latency and video running time.

TABLE 16.1
Download Time and Successful Requests for Video Files with Different Running Times

	Running Time (Approx. Seconds)				
	19	32	58	100	250
Download Time (Seconds)	11	25	50	90	241
Successful Requests (10 users)	10	10	10	2	0

length of a video determines the file size. Together these graphs show that the "longer the film" then the "bigger the file" and that the bigger the file then the longer it will take to download.

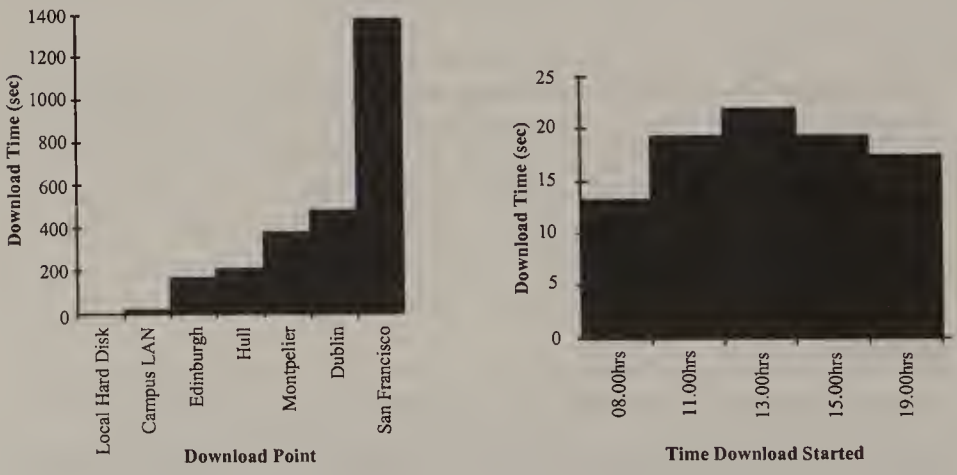
The graphs presented in Fig. 16.4 have a number of practical implications for the design of web pages. For instance, Table 16.1 shows that in one evaluation

only two out of ten casual browsers were actually prepared to wait for 90 seconds while a video file was retrieved from an otherwise unloaded local server. The other eight users all decided to abandon their request before the transfer was complete. None of the users successfully retrieved the file with a download time of over three minutes (Johnson & Kavanagh,1996):

The evidence of this limited evaluation does not imply that all videos should be cut into segments that can be retrieved in under a minute. On the other hand, there are clear advantages to be gained from providing videos with a range of different running times. This enables users to start with relatively small clips and then only download longer films once they are satisfied with the content and production quality of the shorter videos. For the designer, this suggests that greatest attention ought to be paid to shorter clips so that users are 'persuaded' to access the longer material!

Principle 5: Use Mirror Sites if Necessary

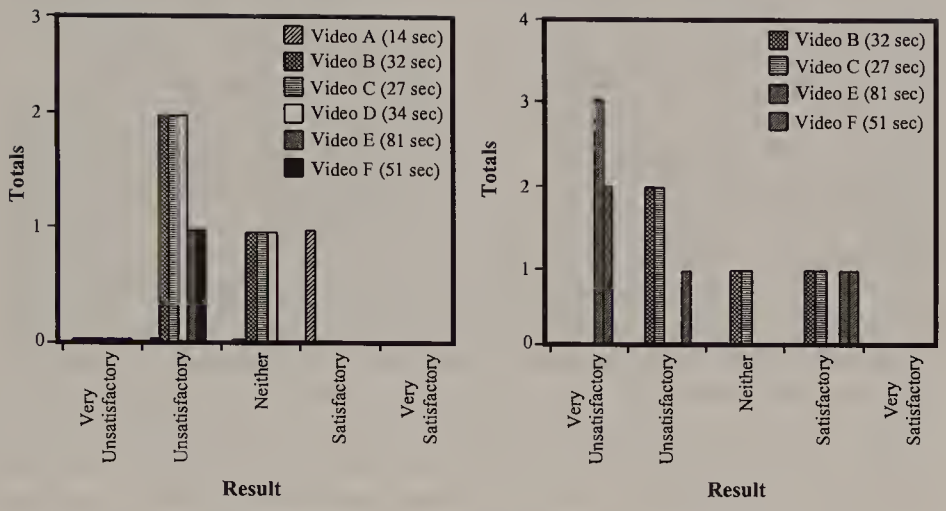
The times shown in Fig. 16.4 were obtained over a local area network with only a light loading. The impact of file size on retrieval delays is proportionately greater over remote networks. Figure 16.5 (a) illustrates the importance of geographical location for retrieval delays. A number of users were asked to download the same video file at approximately the same time of day (Johnson & Kavanagh, 1996). The file was 770 K in size and had a running time of 11 seconds. As can be seen from the graph, retrieval delays are heavily influenced by the user's location. It is important to note that this does *not* mean that delays are determined by the user's geographical distance from a resource. It takes longer for a user in Dublin, which is 200 miles away from Glasgow, than it does for a user in Montpellier, 900 miles away. This is due to a 256 kbit line between Ireland and the United Kingdom. If



(a) Download latency by location

(b) Download latency by time

FIG. 16.5. The relationship between download latency, geography and time of day.



(a) Browsing User (b) Task User

FIG. 16.6. Attitudes to retrieval delays for casual browsers and task-directed users.

designers were providing a resource for use by Dubliners, then the web page would have to be designed as if Dublin were geographically further away than Montpellier. This implies that any attempt by designers to define optimal file sizes, and hence acceptable delays, must be indexed by the relative position of the user to their source (Johnson, 1996). Figure 16.5b adds complexity to this analysis because retrieval delays vary throughout the day. Figure 16.5b) accurately mirrors the peaks and troughs of demand on Glasgow University's computing infrastructure.

Having recognized the impact that both time *and* place have upon retrieval delays for video resources, it is important to ask *what can be done to combat these problems*? Firstly, designers must identify the boundaries of acceptable and unacceptable delays for the users of their resources (Johnson, 1996). These boundaries must be indexed by the geographical location and network resources of the target population. For example, the Hunterian Museum's video archive is intended to support schoolchildren in the West of Scotland. Designers should not, therefore, produce a web page that can only be accessed over the high speed networks that connect Glasgow to Southern California. Secondly, if average delays extend beyond these limits then designers must take steps to reduce the impact of physical, or rather network, distance. For instance, mirror sites can be provided. Video segments can be shortened to reduce file sizes and hence retrieval times. Other delivery mechanisms, such as CD-ROMs, might be considered. Alternatively, designers might argue for improvements to the supporting infrastructure. For example, the delays between Glasgow and Dublin, shown in

Fig. 16.5, may be reduced as a side-effect of a new 512 kbit line across the Atlantic.

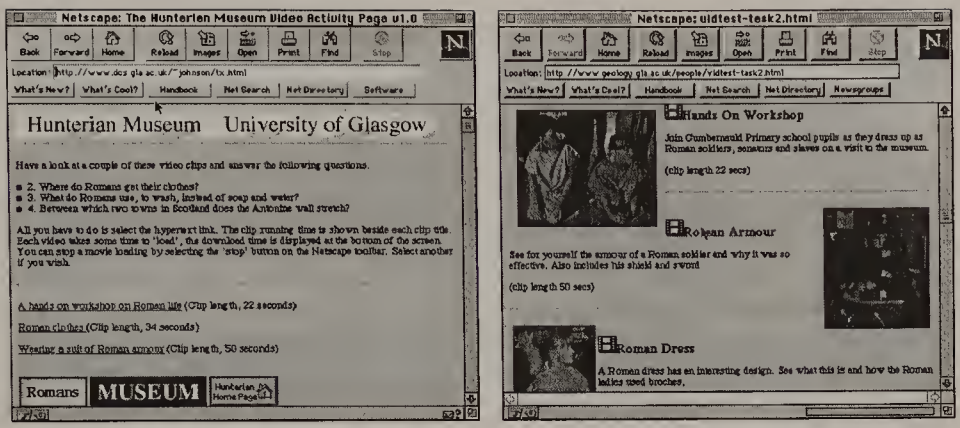
Principle 6: Cater for Differences Between Browsers and Directed Users

Browsers are less tolerant of retrieval delays than people with a clear 'task' (Johnson, 1996; Pjetersen, 1989). Such findings have important consequences for the provision of video material over the web. Designers might support task-directed users by providing longer clips and larger file sizes because these users are generally more tolerant of retrieval delays. However, things are seldom this clear cut. For example, the Hunterian video pages were designed to support school projects about the Romans in Scotland. Exercises were included so that teachers could direct the childrens' search within the video 'archive'. Figure 16.6 presents different attitudes towards retrieval delays. It compares the responses of children who were answering a set of comprehension questions to those who were not (Johnson, 1996). Each group contained the same number of children but two more videos were successfully retrieved in the task directed group.

The results revealed generally negative attitudes to the download time. We could not, however, establish the hypothesis that casual browsers exhibit different attitudes towards retrieval delays than task directed users (Johnson & Kavanagh, 1996). A number of factors may explain the differences between these findings and those cited above. For example, there may be differences in attitudes to retrieval delays between adults and children. These differences may, in turn, be influenced by the extent of the delays involved by particular tasks. Attitudes may also be affected by the unpredictable nature of delays over the web (Johnson, 1995). With so little research into these problems, it is dangerous for designers to extrapolate beyond the experimental conditions of previous investigations. In such circumstances, there are few short-cuts and no substitutes for first-hand evaluations.

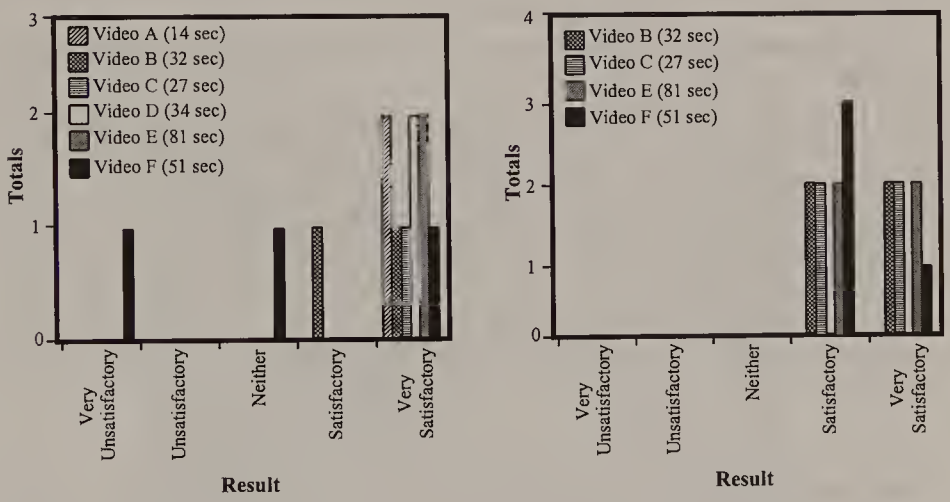
Principle 7: Use Page Layouts to Indicate the Quantity and Quality of Video Material

Principles 3 and 4 argued that designers must optimize the trade-off that exists between production quality and file compression. Such techniques are of little benefit if users cannot recognize the quality of resources from the pages that are used to index them. For example, Fig. 16.7 presents two different approaches to the presentation of video material. In the first page, textual links provide few cues about the content and quality of the video clips. In the second interface, thumb-nail sketches present still images from the videos. This provides a number of advantages. Users can view a small portion of the material without incurring the penalties of a full retrieval. The provision of thumb-nail sketches helps users to plan their subsequent retrieval strategy while waiting for another video to download. The still images also indicate the production quality, or more specifically the frame size, that is used in the digitization of a video. This would,



(a) Textual layout (b) Cinematic Layout

FIG. 16.7. Page layouts for video resources.



(a) Browsing User (b) Task User

FIG. 16.8. Attitudes to retrieval delays for the cinematic page layout.

of course, only work if the designer did not specifically enhance the resolution of the thumb-nail sketch as a means of enticing the user to the video!

Page layouts not only affect the aesthetics of an interface. They can also affect attitudes to retrieval delays. For example, the experiment in the previous section was repeated using the cinematic interface of Fig. 16.7 (Johnson & Kavanagh, 1996). As before, the intention was to identify any differences between casual browsers and 'task-motivated' users. The results of this analysis are presented in Fig. 16.8. The children were more positive about the retrieval delays with the cinematic layout, irrespective of whether they had a task to perform. As

before, however, there were no significant differences between task motivated users and casual browsers.

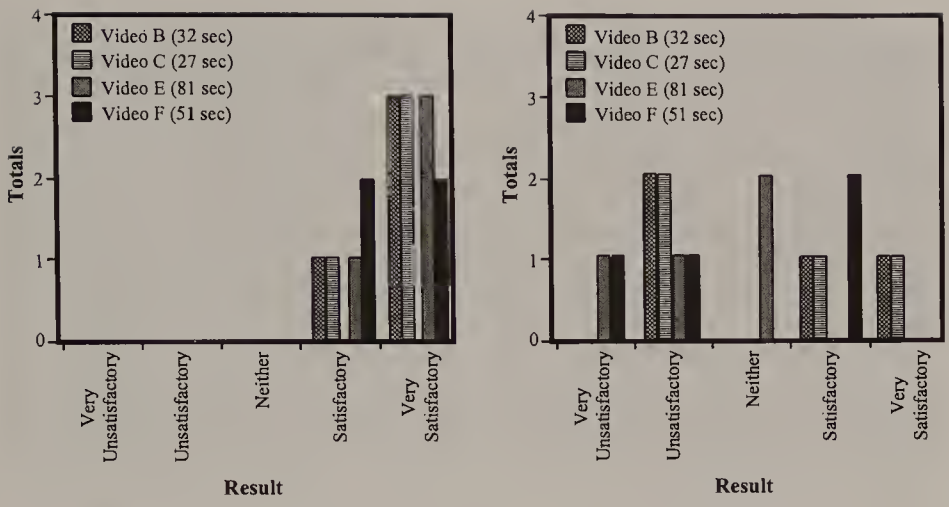
Principle 8: Assess the Impact of Retrieval Delays On Overall Task Satisfaction

Many of the principles in this chapter are intended to reduce the problems that arise when videos take many minutes to download from remote sites. It is important to emphasize, however, that many other factors must be considered when assessing the costs and benefits of information retrieval. For example, task satisfaction may still be high in spite of specific complaints over prolonged delays. This analysis is confirmed by the results of the evaluation described in the previous section. The two groups of "task-motivated" children were asked whether or not they enjoyed their use of the textual and cinematic page layouts (Johnson & Kavanagh, 1996). Their answers are summarized in Fig. 16.9. The qualitative results show very positive attitudes to the cinematic approach. All users were either satisfied or very satisfied even for video clips of up to 81 seconds. Attitudes towards the textual interface were less positive.

It is important to remember that the retrieval delays were constant between the clips in both of the cases shown in Fig. 16.9. In other words, there must be some other influence upon user satisfaction with these different page layouts. We are currently analyzing additional data that was gathered during our evaluation. This data suggests that task users show a more positive attitude to the download but only when the task completion is regarded as successful. When the user does not perform the task in a satisfactory time, i.e. inside the time allowed by a teacher, then the user's attitude is considerably more negative. In the case of the cinematic interface, users were more likely to complete the task because they could more accurately predict which video contained the information that they needed from the still image that was used to index the clip.

Principle 9: Never Underestimate the Element of Surprise

Given that retrieval delays are not the only factor that affects user attitudes towards web-based video archives, it is important to ask *if designers can exploit any other techniques to increase user satisfaction?* One means of achieving this is to introduce elements of surprise or game playing into an interface. These approaches are especially appropriate in educational software where video clips can be used as a reward for good performance. They can be used to retain user interest. They can also increase levels of acceptance for material that is presented in other formats (Neal, 1990). Figure 16.10 shows how an element of surprise can be introduced into the delivery of video material over the web. This interface to another section of the Hunterian Museum's web site presents an image map of the route taken by Captain Cook in his third voyage of discovery (1776–1780). By selecting certain different areas of the screen, users can access information about



(a) Cinematic Interface (b) Textual Interface

FIG. 16.9. Task "satisfaction" for the cinematic and textual page designs.

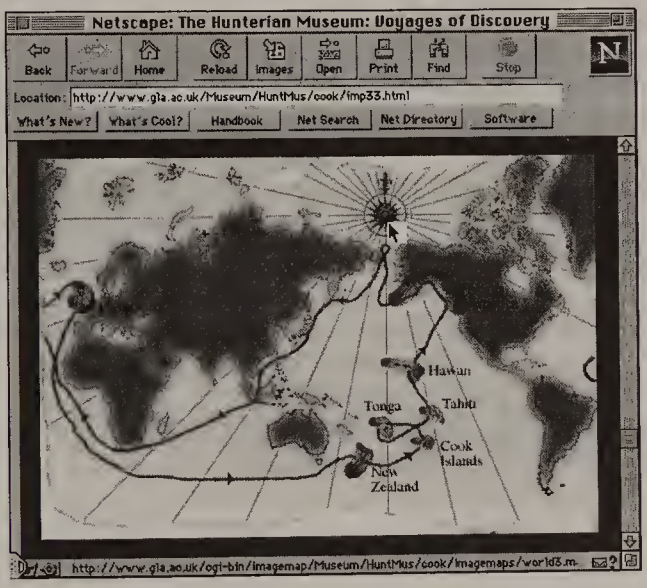


FIG. 16.10. Image map showing Captain Cook's second voyage of discovery.

the areas that Cook visited during this voyage. Some selections result in question and answer sessions, others provide access to textual and photographic records, further areas provide access to video footage. The intention is to support 'learning by exploration' within the overall topic framework provided by the map (Carroll, et al., 1985).

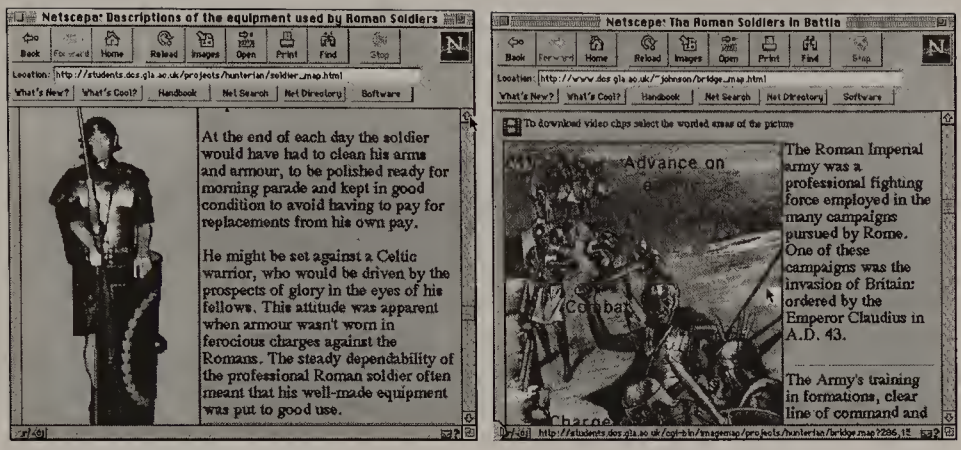
The interface in Fig. 16.10 violates some of the principles that were advocated in previous sections. For example, the cinematic layout of Fig. 16.6 was successful because it explicitly represented the content and production quality of each video. This information is deliberately hidden in Fig. 16.10 because the user cannot even tell if there is any video material linked to the page. This illustrates an important point about the application of guidelines. Designers should *not* attempt to achieve all of the principles that are advocated in this chapter. They *must*, however, have a clear reason for violating any guidelines that they do reject. For example, the web page shown in Fig. 16.10 violates our principles in favour of preserving the principle of *direct manipulation.*

Principle 10: Never Take A Designer's Word For It

This principle is included as a warning against introspection. The previous section argued that designers must only violate a principle if they have a clear reason for doing so. This section argues that designers must seek empirical evidence to determine whether these reasons are justified within the context of a particular interface. For example, Fig. 16.11 presents two different techniques that can be used to prompt users about the location of video material within an image map. These layouts do not achieve the element of surprise that was intended in Fig. 16.10. The intention is, however, to index video clips with information about the context of particular objects or subjects. In other words, a video about the use of a shield can be retrieved by selecting a shield. A video about the roman tactics during an advance can be indexed by selecting a picture of an advance. The interface on the left indicates the presence of a video clip by outlining areas of the image in a 'special' color. In this interface, it was a bright purple. The spear and shield in Fig. 16.11 are outlined in this manner. The layout on the right uses an image map to associate video clips with textual labels. Combat, Advance One Enemy and Charge all index different clips. A third design was also produced in which colored spots were placed next to areas on the image map that could be selected to retrieve a video. These were similar to the pins that might be placed into a postcard or map.

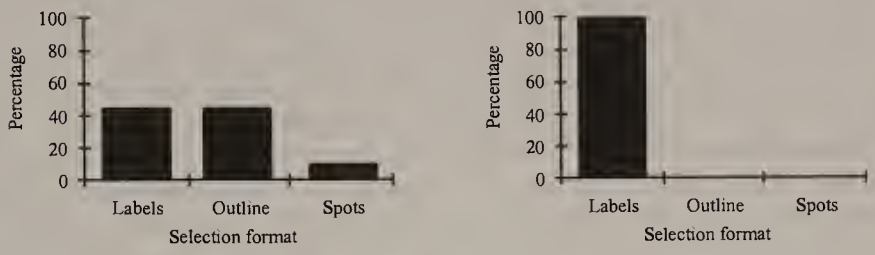
An evaluation was conducted to compare the users' reactions to the designers' expectations for these different layouts. Thirty school-children from the West of Scotland were asked to select *which of the interfaces shows the selectable areas most clearly?* Web-page designers from the staff and undergraduates of Glasgow University were then asked to predict which of the two interfaces would best present the selectable areas for our target population. Everyone in this second group had considerable experience of web design and had undergone at least two years of training in Human–Computer Interaction. Figure 16.12 presents the results.

The findings came as a considerable surprise to the development team. Nobody had predicted the apparent unanimity amongst the potential users of the system. Many of the designers had even considered abandoning the labeled presentation technique in favor of the outlined approach. The important point, however, is not that users preferred one layout over the others. Nor is it that



(a) Outlined Areas (b) Textual Labels

FIG. 16.11. Image map interfaces to video pages.



(a) Designer's Reaction (b) User's Response

FIG. 16.12. Differences between the designers and users of web-based video systems.

designers cannot accurately predict user preferences. It is, rather, that there is no substitute for direct user trials given the lack of any theoretical framework or applicable empirical research that might otherwise guide the development of web-based video archives.

SUMMARY

This chapter has presented ten principles that support the provision of video material over the World Wide Web:

Principle 1: Avoid Video for Video's Sake
Principle 2: Make Sure That Users Can Access the Video
Principle 3: Maximize the Production Quality

Principle 4: Edit The Clip to Minimize File Size
Principle 5: Use Mirror Sites if Necessary
Principle 6: Support Both Browsers and Directed Users
Principle 7: Indicate The Quantity and Quality of Video Material
Principle 8: Assess Overall Task Satisfaction
Principle 9: Never Underestimate the Element of Surprise
Principle 10: Never Take A Designer's Word For It...

It is important to stress that these are heuristic guidelines. They have been compiled during the development of existing web sites. It, therefore, follows that other practitioners may identify omissions from our list. For example, we have not provided any direct guidance on the direction of video footage. Further work intends to address this limitation by further investigating the use of "sharp cutting" techniques in conventional films (Bernstein, 1988). This approach uses a number of cinematic techniques to maximize the impact and information content of relatively short film clips.

Finally, it is important to return to our tenth principle: *never take a designer's word for it....* This emphasizes the need to obtain direct empirical evidence for claims about the 'usability' of web pages. Unfortunately, many of the methodologies that have been developed to examine "conventional" interfaces cannot easily be applied to evaluate distributed, information retrieval systems. The traditional measures of precision and recall completely ignore the temporal problems that arise during interaction with the web (Teal and Rudnicky,1992). Without further work in this area, it is important that designers are provided with principles or guidelines that build upon the successes and failures of previous interfaces. Without such guidelines, users will continue to tie up their machines for long periods while mediocre footage is retrieved across heavily loaded networks.

ACKNOWLEDGMENTS

The web pages described in this paper were developed with Robert Anderson, Ian Dunn, Sean Hundtofte, Thomas Mackenzie and John Montgomery of the Department of Computer Science, University of Glasgow, and Jim Kavanagh, University of Paisley. The video clips were produced by this team and by the Scottish Television "Skoosh" crew with Lindsay Hill (the presenter) and Beverley Morrison (producer). Video Stills © Skoosh, Cool Crew Productions for Scottish Television. Special thanks are due to Jim Devine and the staff of The Hunterian Museum for their help and encouragement in this work.

REFERENCES

Bernstein, S. (1988). *The technique of film production*. London: Focal Books.

Carroll, J. M., Mack, R. L., Lewis, C. H., Grischkowsky, N. L., & Robertson, S. R. (1985). Exploring a word processor. *Journal of Human Computer Interaction, 1*(1), 283–307.

Dix, A. J. (1987). The myth of the infinitely fast machine. In D. Diaper & R. Winder (Eds.), *People and computers III* (pp. 215–228). New York: Cambridge University Press.

Gray, P., England, D., & McGowan, S. (1994). XUAN: Enhancing UAN to capture temporal relationships among actions. In G. Cockton, S. W. Draper, & G. R. S. Weir (Eds.), *People and computers IX, Proceedings of Human Computer Interaction '94* (pp. 301–312). Cambridge, England: Cambridge University Press.

Johnson, C. (1993). A formal approach to the presentation of CSCW systems. In J. L. Alty, D. Diaper, & S. Guest (Eds.), *People and computers VIII, Proceedings of Human Computer Interaction '93,* (pp. 335–352). Cambridge, England: Cambridge University Press.

Johnson, C. (1995). Time and the web: Representing and reasoning about temporal properties of interaction with distributed systems. In M. Kirby & D. Diaper (Eds.), *Proceedings of Human Computer Interaction' 95* (pp. 39–50). Cambridge, England: Cambridge University Press.

Johnson, C. (1996). Electronic gridlock, information saturation and the unpredictability of information retrieval over the WWW. In F. Paterno & P. Palanque (Eds.), *The application of formal methods*. Berlin: Springer-Verlag.

Johnson, C., & Kavanagh, J. (1996). Electronic gridlock: The problem of download latency for video on the web. In A. Blanford & H. Thimbleby (Eds.), *Adjunct Proceedings of Human Computer Interaction 96* (pp. 78–89). London: British Computer Society.

Kuhmann, W. (1989). The stress inducing properties of system response times. *Ergonomics*, *32*(3), 271–280.

May, J., & Barnard, P. (1995). Cinematography and interface design. In K. Nordby, P. Helmersen, D. Gilmore, & S. Arnesen (Eds.), *Interact '95* (pp. 26–31), London: Chapman and Hall.

Neal, L. (1990). Implications of computer games for systems design. In D. Diaper, D. Gimore, G. Cockton, & B. Shackel (Eds.), *Interact'90* (pp. 93–99). Elsevier.

Pjetersen, A. M. (1989). A library system for information retrieval based on a cognitive task analysis and supported by an icon based interface. In *ACM SIGIR Proceedings*. Boston: ACM.

Teal, S. L., & Rudnicky, A. I. (1992). A performance model of system delay and user strategy selection. In P. Bauersfeld, J. Bennet, & G. Lynch (Eds.), *Proceedings of Computer Human Interaction'92: Human Factors in computing systems* (pp. 295–305). New York: ACM.

Part V: COLLABORATION AND VISUALIZATION

17 Graphics Design on the Web

Eric N. Wiebe
Julie E. Howe
North Carolina State University

One of the appealing elements of the World Wide Web (WWW or Web) lies in its ability to integrate, into a single document, both formatted text and graphics. The hypermedia capabilities of the Web add tremendous flexibility in the creation and communication of information. With flexibility, however, come ever-increasing possibilities of using techniques and media that are both computationally inefficient and poor practice from a human performance standpoint. A review of current information found on the Web reveals thorough and thoughtful design guidelines which discuss issues concerning both text and graphics (e.g., Berners-Lee, 1995; Levine, 1995; Lynch, 1996) Most of these documents, however, concern themselves more with issues of aesthetics and computational efficiency than with human factors elements. This chapter provides an overview of some overarching human factors issues involved with the design and use of graphics on the Web. These larger theoretical points are also paired with a discussion of how they might apply to specific application areas on the Web. The goal is to go beyond strictly computational performance and aesthetic issues and look at issues pertaining directly to visual performance: how to improve the speed, accuracy, and usability of communicating graphic information.

As with printed documents, it is difficult to divorce graphics from textual elements when evaluating an online document. Inevitably, there is a high degree of interaction between text and graphics so that many measures of task performance depend on the design and placement of both elements. It is often hard to classify elements as being purely textual or purely graphic; a graphic may contain text in it, or a block of text may be formatted using color, borders, or other "graphical" elements. But, the complexity and scope of the issues involved in Web document design argue for some division when attempting a discussion of performance issues. A good starting point in designing a document for the Web may be to ask some simple questions (Slack & Conati, 1992): First, what categories of to-be-communicated information can be encoded graphically? Second, when are graphics the most "effective" medium of presentation for a given piece of information? For example, would a graphic showing the hierarchical structure of documents in a Web site be more effective than a bulleted text outline? Finally, what is the most "effective" graphical encoding of

information? What colors, shape, textures might be most effective for segregating the different categories of information? This chapter, for the most part, focuses on the issues raised by this last question.

RESEARCH ISSUES ON PERCEPTUAL ELEMENTS

An understanding of human visual perception can be a powerful tool for designers of online information systems such as the Web. Roughly speaking, perceptual processes are those mental processes (in this case, those involving vision) that are essentially effortless or automatic. One approach to perception is to evaluate how different physical characteristics are processed. These physical characteristics can be subdivided into graphic features such as color, location, texture, orientation, shape, and size (Baker & Wickens, 1995; Haber & Wilkenson, 1982). Such graphic features are particularly important to the design of information sources because of their associative, selective, and ordering capabilities. Poor use of these features can add greatly to the mental effort needed to understand and extract required information from the display, whereas deft use of graphic features can clarify the organization of information on a Web page through the use of effortless perceptual processes.

No matter what the visual medium, color is one of the most common graphic features used to organize information. Numerous comprehensive writings specifically concerned with human factors applications of color displays show how proper use of color can make marked differences in the usability of a display (c.f. Christ, 1975; Silverstein, 1987; Thorell & Smith, 1990). Color's effectiveness comes from the flexibility with which it can be used to organize, order, or highlight. Most beneficially, it can be used to code information without taking up screen "real estate"; that is, it is a nonspatial cue used to alter the characteristics of existing spatial elements.

Texture also has potential as a method for coding information. Even though texture is a more complex building block than color—patterns contain multiple basic features such as orientation and size—there is reason to believe that some elements of texture are perceived preattentively (Julesz, 1981). Unfortunately, the lack of robust classification schemes limits textures' ability to be specified and therefore its usefulness as a precision design element (Rao & Lohse, 1993). In addition, texture, as a spatial factor, is sensitive to the scale at which it is displayed. The limited resolution of computer displays commonly used for Web browsing, coupled with constrained screen size, limits the pattern size and detail at both the upper and lower extremes.

Arrangements of graphic primitives such as lines, areas, and volumes have configurational properties that cause them to be considered perceptually as part of a whole rather than as individual elements (Garner, 1981). In this case, features can be ascribed both to the individual primitives and the object (the constituted whole). For example, the lines that make up an arrow each have an orientation, but the completed arrow has its own unique orientation. When creating Web

graphics, designers must be concerned not only with the properties of the low level elements such as individual lines, but also the sometimes unique properties of higher level objects that arise from arrangements of these primitive elements.

The process of perceptual organization can lead to the perception of both two-dimensional and three-dimensional form (Hochberg, 1981; Zusne, 1970). Because the human perceptual system has evolved in a 3-D environment, there is a bias toward trying to make sense of a pattern as a three-dimensional form. As computer displays lack many cues found in a true 3-D environment, viewers often make 2-D interpretations of many patterns, intended or not. Animation is one way of enhancing the 3-D perception of form. Another way is to use so-called virtual reality tools that allow for the addition of interactive capabilities, adding even more perceptual cues. The advantages of 2-D versus 3-D representations must be evaluated from both a computational standpoint and a human factors perspective; designers must ask, "In view of the display capabilities of current Web browsers and the task at hand, what will be the most effective display method?"

Cognitive Elements

One way to distinguish between perceptual and cognitive processes is to use an information-processing model based on limited resources (Baecker, 1987). Under the assumption that most perceptual processes are automatic and effortless (i.e., not consuming resources), the processes that need mental resources can be considered cognitive. An important operational impact of this model is that there is a tradeoff between resources consumed and quality of information presented. In addition, there is also a tradeoff between resources consumed and parsimony of displayed information structure and existing mental structures. These operational rules of thumb become clearer as specific examples are discussed.

Graphics, when properly designed and implemented, can be powerful tools for information encoding (Alty, 1991; Haber & Wilkenson, 1982; Hodes, 1994). Graphics can increase the size of the information chunk to be processed: The same resources used to process a single sentence of a paragraph can be used to process a graphic that summarizes the entire paragraph. In many cases, the situation may not be either-or; graphics may be simply the only effective way of expressing a concept. When a user's knowledge in an area is growing, graphics can be powerful synthesizers. Graphics can be used to link, summarize, and coordinate information with minimal resources. Often, the key is to identify the essence of the information to be conveyed and to find ways of embedding the information in a graphic design. The problem is pinning down cognitive factors such as communicativeness and design quality to develop a meaningful measure of effectiveness (Lin, 1995).

Another way in which mental resources are consumed is in the recall of previously learned information. Users of computer software acquire knowledge about both the functionality of the software and the information revealed through the interface as they use the software. This knowledge can range from memorizing a simple sequence of overt actions to a mental model composed of metaknowledge of system or information structure (Carroll & Olson, 1990). An

example is the difference between knowing that a particular icon on a page takes the user back to the home page where you started versus understanding enough about the overall structure of a site to be able to click efficiently through links (never visited before) to a piece of information. A review of literature on mental models and knowledge acquisition in the human computer interface (HCI) field shows that many issues pertaining to typical software interfaces are not important in Web design. For example, on a typical browser, the commands are not very hard to learn, and there is little need to know how the system (i.e., WWW servers, Internet protocols or client architectures) works to extract information from the Web. Like database software, though, the structure of information available through the interface is extraordinarily complex. It is central to look at how graphics can support the development of appropriate mental models or tap effectively into existing models to support use of information repositories on the Web.

Task–Design Interaction

All evaluations of the interface are, at some level, dependent on the task performed (Frey, Rouse, & Garris, 1992; Sobol & Klein, 1989). When assessing perceptual or cognitive effectiveness of an interface in an operational environment, it is always critical to carefully define the task a user performs with the interface. To add to this complexity, the evaluation is also dependent on the background and experience of an individual user (Cherry, Turner, & Rockwell 1992). Predicting a user's task and experience level can be extremely problematic when it comes to the Web: So many sites are visited by a widely diverse audience with a wide range of task motivations and experiences. For those designing intranet applications for specific tasks or developing sites dispensing information to a well-defined population, the task is simpler. For other sites, only the most general assumptions about user population may be possible. Still, fundamental principles of perception and cognition can be applied to a design to provide an overall enhancement to usability.

Web Issues

Designing a user interface, such as a Web page, requires balancing a user-centered approach against the need to recognize both the strengths and weaknesses of the available technology. The following graphic interface design goals were set for the Alto/Star (Xerox Corporation, http:/www.xerox.com), an early object-oriented graphic user interface: to match the graphics with the medium (in this case a Web browser on a personal computer); to strive for appropriate effect in how the interface responds to users; to develop a consistent vocabulary of (graphic) interface elements; to create visual order and user focus; to create a revealed structure so that the system architecture and functionality are as readily apparent to users as possible; and, to create the illusion of manipulable objects in the interface (Verplank, 1990). Even though these goals were set for an interface developed in the late 1970s (a millennium in the microcomputer world), they still hold true for the development of Web interfaces. These design goals can be

summed up in a single statement: Graphic elements, like text, should be enlightening, not distracting.

Common Web Graphic Techniques

Many common graphic techniques used in Web page design can be enhanced through thoughtful application of basic cognitive and perceptual principles. One basic principle is that humans have the capability to digest information through multiple modalities. In many cases, redundant information channels can enhance the usability of a Web page. For example, Hyper Text Markup Language (HTML) provides a way of attaching text to a graphic so that if users prefer not to display graphics, they can see a short text summary or title for the graphic. Similarly, text can be embedded in a graphic as part of the bitmapped image, to provide a redundant or complementary information source (e.g., the graphic representation of the front view of a house is used extensively in hypermedia documents as representing a link to the "home" page). The text "home" under the house can be designed so that it does not distract from the graphic effectiveness of the house while it provides an alternate cue for those who may not recognize the symbol. When graphics are used as a hot link, a redundant text link can also be provided outside the graphic. For a graphic formatted as an image map with multiple links, a series of text links separated by a salient character (e.g., the pipe symbol, |) can be placed under the graphic as an alternative to navigating through the graphic. This tactic is particularly important if the meaning or location of the links in the graphic is at all unclear.

One way of clarifying that a graphic is a link to another location (thus heightening its importance on the page), is to outline it with a border. A border, by default, is placed around all linked graphics and can also be placed in a graphic to identify links embedded in a portion of an image. The default border can be removed, to leave the color contrast of the edge of the graphic against the background to define a border. Even more extreme is the transparency option in the GIF89a (Graphics Interchange Format) standard (Compuserve, 1990) which blends the graphic even more completely into the background. These borders act not only as cues to the functionality of the graphic, but also as a means of focusing attention. Work by Brown (1991) has shown that borders can shorten search time in word displays. As with all design decisions, designers must keep task goals in mind and weigh the pros and cons of this cue: Borders can take attention away from other portions of the page, and can also affect the processing of information inside themselves (Swierenga, Boff, & Donovon, 1991). Therefore, the sharpness of the color contour between graphic and background and the regularity of the border shape can both be modulated to determine the degree to which a designer attracts attention to the region and to the edge of the graphic itself.

In addition to borders, another way to enhance the contrast between discrete graphics and background is to manipulate the color, pattern, or both, of the background. With careful planning, this technique can be a powerful technique for coding different documents or groups of documents in a Web site.

When used for solid backgrounds, colors need to be chosen with adequate contrast to text and graphic elements on the page to give optimal performance and comfort (Thorell & Smith, 1990). Contrast can be achieved by manipulating either the lightness component or the hue-saturation component of the color. Medium-gray is often used as a default background simply because its lack of hue-saturation puts it at the extreme of one scale and very light or very dark text also contrasts on the lightness scale. In addition, low saturation colors recede, allowing foreground elements in higher saturation colors such as icons, text, and titling, to be more apparent (Truckenbrod, 1981).

Background patterns should be used with extreme caution. In addition to adding overall visual complexity to a display, contrasting elements in the pattern can selectively obscure or enhance other visual elements in a display, with unintended results. Such a pattern, as another graphic element in a display, will compete for attentional resources and may lead to lower performance on searching and other tasks (Zwaga & Duijnhower, 1995). Unfortunately, the same technique that reduces competition with foreground text and graphics (i.e., low-saturation colors and low contrast-low frequency patterns) also makes the background hard to discern. Often it comes down to whether a designer wants the focal point of attention to be in the foreground or background.

Creating a blinking element is another way of drawing attention. Again, task goals should guide decisions about the usefulness of a flashing element (Verplank, 1990). If there is a strong desire to draw attentional resources away from other elements in the display, flashing elements achieve this goal. A compromise may be to use various animation techniques to create a dynamic image sequence for a short period. The GIF89a standard allows an image file to actually consist of multiple frames displayed in a streaming fashion in a Web document. This stroboscopic technique can be effective to display simple animations. Although they can be set to loop indefinitely, such elements can be set to play through once, long enough to initially get users' attention, but short enough not to compete with the rest of the page. More sophisticated programming tools such as Java (Sun Microsystems, Corporation http://www.sun.com) and Shockwave (Macromedia Corp., http://www.macromedia.com), can be used to create more sophisticated effects, including animations that stop when clicked on. The cardinal rule is that blinking and other animated elements should be used sparingly and only when they serve a specific function on a page.

Scientific and Technical Data Representation

True to the historical roots of the Internet, the Web has been extensively used as a means of exchanging and publishing scientific and technical information. As is true for print media, researchers' work (Cleveland, 1985; Gillan & Lewis, 1994; Kosslyn, 1989) linking perceptual and cognitive theory to the design of charts and graphs is equally pertinent to their display on the Web. Morse (1979) established two primary goals for the display of data on the computer. First, the perceptual effect should be equal to the magnitude of the datum and at the degree of accuracy needed by a user. If a user must know specific data values, for example, these

should be redundantly provided as alphanumeric values. The second goal is to design the display so that the least (mental) effort is needed to assimilate the information. In short, the power of graphics (when used properly) lies in their ability to integrate and summarize. Both human factors research and operational experience of graphic display designers are important in assessing this dimension. One methodology emphasized by Tufte (1983) is the determination that all graphic elements in a display are closely tied to the representation of a data dimension. Human factors research can help determine when the addition of more graphic marks no longer enlightens, but rather distracts from the display of data.

As mentioned earlier, color is a powerful tool for reducing search times and organizing and grouping data. The effectiveness of color coding is most salient in cluttered or complex displays, but it is also measurable even on relatively simple designs (Christ, 1975; Hughes & Creed, 1994). Color as a redundant code also shows advantages over shape-only coding (Brown, 1991; Hughes & Creed, 1994; Zwaga & Duijnhower, 1995), even with practice (Jubis, 1990). Research has shown that color is processed separately but in parallel with spatial properties such as orientation, shape, and direction of movement (Treisman & Gelade, 1980). Whether in conjunction with shape or by itself, color is effective at enhancing visual performance, but only when it codes a data dimension relevant to the task.

The type of scale that data represent is an important consideration when choosing a coding strategy (Bertoline, Wiebe, Miller, & Nasman, 1995). Typically, hue is used as a strategy for coding all types of data dimensions but, in fact, it is unsuitable for a common type: quantitative data. Although there are common conventions for the use of some hues (e.g., red is hot and blue is cold), there is no way of creating equal perceptual distances between hues; the "distance" between red and blue cannot be described in units (Meyer & Greenberg, 1980; Thorell & Smith, 1990). Merely because a graphics software package offers color palette tools that display Hue, Saturation and Lightness as digital values, there is no way to guarantee a monotonic relation to perceptual responses to these colors. Perceptually, uniform color distance is extremely hard to calculate and is dependent on many factors, not the least of which is the computer operating system. Apart from perceptual response, only a few "safe" colors (216) can be used in GIF images that display similarly on Macintosh and Windows-based 8-bit displays. For images where color management is critical, JPEG (Joint Photographic Experts Group) images allow for more color control and an expanded palette. Regardless of image format, with all color coding of data dimensions, legends are obviously an important design element to equate specific displayed colors with data values.

Research by Wickens and colleagues (Wickens & Andre, 1990; Wickens & Carswell, 1995) has shed light on the linkage between perceptual proximity and task. Perceptual closeness can arise not only from spatial dimensions such as orientation and physical proximity, but also from nonspatial factors such as color. More specifically, there is parsimony in using perceptually separate data displays for tasks that require focused attention on individual values; perceptually

integrated tasks favor the use of integrated displays, where the graphic elements form unified objects. Besides spatial separation, perceptually separate displays can also be created by using color codes to segregate values. Color separation of elements then allows users to effectively focus on individual values in a display. On the basis of its use to create perceptual separation, color coding for the purpose of segregation in an object display would seem to be counterproductive to integrative tasks. The effect of color, however, may come down to what information source is supporting the integrative task. If spatial properties of an object, including higher order, or emergent, features such as parallelism support an integrative task, the nonspatial color dimension may have only a marginal effect on performance.

By classifying features as to their degree of homogeneity, predictions can be made about the likelihood of their binding in such a way as to make it easy to interpret the grouping as a whole (as in an integrative task) but difficult to interpret them separately in a focused attention task. Using feature attributes with separate perceptual analyzers is one way to keep features from binding. The most common examples of separate perceptual feature attributes are shape and color. By using separate perceptual channels, a designer can have the spatial and nonspatial (i.e., color) codings be redundant or be used to code different data dimensions. Although there is disagreement as to whether the interaction of these channels is truly negligible, there is support of the general utility of this approach (Baker & Wickens, 1995; Treisman & Gelade, 1980). The important message is that information can come from many sources; these sources must be identified, and they must be matched to the task at hand. Some of these features, such as emergent features, may not be directly represented in the graphic, yet they may have unintended or unwanted influence over how the data is interpreted.

In addition to these examples of 2-D shapes representing data dimensions, 3-D objects can also be used. These objects, however, bring a unique set of problems in mapping data dimensions to spatial dimensions of an object. Three-dimensional objects, often referred to as glyphs, commonly use metric (size) and shape changes as a way of mapping data, yet these two methods support different levels of accuracy in interpreting data values (Wiebe & Converse, 1996). Glyphs usually have shape constancy whatever viewpoint the objects are projected onto a screen from, as long as the shapes' critical features are not occluded. Accurate interpretation of metric variation, on the other hand, interacts with the viewpoint and leads to worse performance when the varied dimension of an object is foreshortened. Size variation is a basic method of coding data in 2-D graphics such as bar graphs, partly because it is such a direct method of encoding quantitative variables. Carrying this technique to three dimensions must be done carefully. Going to a nonspatial dimension such as color, unfortunately, is not always a helpful alternative. Lightness, an effective color component used for mapping quantitative data in 2-D graphics, is typically used to represent a light source on the 3-D object, a practice that can lead to a conflict between support of objectness and the mapping of data values.

Virtual Reality Techniques

The representation of 3-D objects and environments on computers, commonly found under the head of virtual reality (VR), has been a point of focus in the computing community for some time and has certainly carried over to the Web. Virtual reality markup language (VRML) is an industry sponsored standard for the specification of virtual environments on the Web and other interfaces, with the recently released 2.0 version being widely accepted in the computing and entertainment industries (VRML Consortium, 1996). Although great strides have been made in the representation of virtual environments, a computer-generated world is still an impoverished visual environment. Because of the limited processing power of a typical computer connected to the Web coupled with the bandwidth constraints of the Web, the possibility of high-fidelity virtual worlds is still distant. On the basis of accepted criteria for simulator imaging, the Web, unfortunately, lacks many capabilities currently available in high-end simulators (Padmos & Milders, 1992).

Ellis (1995) defined three levels of virtualization in a VR environment: virtual space, virtual images, and virtual environments. Virtual space is the process by which a user perceives a 3-D layout of space when viewing a flat surface with pictorial cues in space: perspective, shading, occlusion, texture gradients. This process is akin to picture viewing and—although not photorealistic—is largely possible on the Web. Virtual imaging occurs when oculomotor cues from accommodation, vergence, and (optionally) stereopsis are present, although not necessarily consistent. These cues largely depend on specialized equipment currently not available to most people using the Web. Some autostereoscopic techniques (i.e., no viewing apparatus needed) are usable with computers, but their effectiveness in depth discrimination is not clear (Adelson, Allen, Badre, & Hodges, 1993). In Ellis's definition of the third level of virtualization, virtual environments, interactive elements, such as observer-slaved motion parallax and wide field of view, are added. Some interactive elements are available in VRML-specified worlds, but are typically controlled with a mouse. Wide fields of view, as previously mentioned, are not yet available on mainstream computer displays.

In addition to the monoscopic pictorial cues offered by static images, interactive, dynamic movement through a virtual world adds considerably to the salience of objects and spatial layout in the world. Dynamic imaging can help remove many ambiguities that result from viewing the scene from a considerably impoverished source (Haber & Wilkenson, 1982; Kaiser & Proffitt, 1989). The perceptual organization of 3-D form through changing viewpoint of an object (often referred to as the kinetic depth effect, or structure from motion) is a way in which the limited computing resources of the Web can be used to maximum benefit. In many cases, the trade-off of realistic movement at the expense of geometric or rendered realism is preferred because it provides more of the type of perceptual information sought from a virtual environment (Baker & Wickens, 1995). When designing an interactive environment, the rate at which a browser can stream the animation off the Internet or display it on a computer is an

important consideration. Research on stroboscopic movement has indicated that both the frame rate and the amount of displacement between frames can be critical factors in designing an effective perception of 3-D form (Petersik, 1980; Wiebe, 1994).

Increasing realism of the 3-D object represented on a computer does not always provide an increase in recognition. Studies have shown that static line drawings of 3-D objects, void of any shading, convey the necessary information for many recognition tasks (Biederman & Ju, 1988; Ryan & Schwartz, 1956). Recognition tasks with both static and dynamically rotating objects have shown lower error rates with line drawings than with shaded models (Wiebe, 1994). Again, both the type of tasks being performed and the complexity and type of geometry probably influence the effectiveness of line drawings over shaded models. When discussing line representations, it is important to note whether occluded surfaces are calculated and removed from view. Wireframe models, such as those used in CAD and VRML type 3-D models, simply represent the edges of the object where surfaces join and do not perform any hidden surface removal. Since occlusion is a critical pictorial cue, this type of model has been shown to benefit from the removal of edges that are hidden from the chosen point of view (Sanford, Barfield, & Foley, 1987).

Graphics Representing Information Structure

One of the most challenging aspects of designing a Web site—or any type of hypermedia application—is to find ways of supporting and developing a user's mental model of how information is structured in a site. Graphics can be a powerful tool for representing both spatial and nonspatial structures in a tangible, compact form (Haber & Wilkenson, 1982; Hodes, 1994). Visualization of information structure has been identified as a key research area in hypermedia design (Waterworth & Chignell, 1989). Examples of prototype interfaces of hyperstructures have ranged from 2-D graphs (Domel, 1995) to 3-D interactive spaces (Apple Computer, 1996).

When using graphics to represent information structures, one issue to be addressed is matching the levels of abstraction and the percentage of total information (aggregation) to the task being performed (Frey et al., 1992). Understanding the task helps determine whether a more succinct, abstract display with the potential of displaying more of the entire information set supports the task as well as a more concrete display that may be able to display only a portion of the entire information set. In each level of abstraction and aggregation, a balance of complexity of the graphic against the completeness of information presented must be achieved. For example, a company's home page may indicate its different divisions with icons or logos and thus maybe provide a more succinct and compact representation of the divisions, which might, in turn, allow detailed information about each division to fit on the page. On the other hand, if a user who accesses the site does not know what the icons represent, the associated information may not have much meaning.

Designers of hypermedia systems also face special usability issues pertaining to navigation (Valdez, Chignell, & Glenn, 1995; Waterworth & Chignell, 1989). Users of the Web must be able to find information efficiently, find their way back, make sure they see relevant information at a site, and not limit their freedom to explore.

Space, per se, is not intrinsically memorable. Although hierarchical list formatting, through the use of HTML tags such as "ol" and "ul", is the most common method of creating a visual structure of the site, discrete graphics and image maps can be used to provide a more flexible method of revealing structure. Graphic techniques such as spatial separation and orientation can be used to order links in a site. Ordering links vertically or horizontally can imply an ordering or grouping of sites. Similarly, boundaries created with linework or color changes can be used to group sites and represent connectivity. Whether a user has visited a site previously or not, it is hard to retain situational awareness while making decisions about how to navigate a site because of the number of possible routes to and from a particular page. Graphic techniques can counter this problem by providing distinct, memorable landmarks with cues easily recalled while a user navigates a site. A common technique is to provide abstract graphic symbols (icons) that are likely to be remembered due to their distinctive perceptual or cognitive properties. An icon might have a memorably unique color or shape, or it might be associated in memory with a previous experience. The use of icons has been shown to help in learning computer databases (Merwin, Dyre, & Humphrey, 1995), a task which has navigation challenges similar to the Web.

Support for an effective cognitive map of a site comes not only from a logical structure of pages and links, but also from supporting graphics that represent this structure. Coding methods such as color, shape, and location, which highlight select information, are one way to alter how an information structure is perceived. These methods implicitly set a hierarchy of information importance for the display. As a standard convention, a border or underlining of a set color is used to indicate links in a Web document. Similarly, a graphic can be designed with a spatial structure that matches a previously learned structure or a standard convention. For example, users are used to looking at the bottom or top of a Web page to find navigation links. Here, a user may find the graphics used on VCR controls, which have become nearly universal symbols for local navigation, even though they do not reveal anything about the overall structure of a site.

A difficult design challenge is finding a graphic that reveals something about the overall structure of a site and is available when a user needs it, yet is compact enough not to take up too much space. The use of frames on a Web page, which allow for a portion of the window to be retained while changing pages, is one way to make sure that navigation links are always available. This technique commits a portion of the window to displaying this information. More sophisticated tools, such as Java, allow for flexible, modal displays of information. If developed properly, these tools can allow better congruence of users' goals and knowledge with what information needs to be displayed. With

current technology, text is very easy to "build" on the fly, but modification of image maps and other graphics presents more of a challenge.

In determining how to code the information structure of a site, the possible representations can be defined as symbolic, spatial filling, or a combination of both (Dumais & Jones, 1985). As previously discussed, spatial coding takes up real estate and there are physical limits to how much information can be represented. Text, as a symbolic method of representing a site, has the advantage of succinctly linking Web pages. There is no evidence that navigation through graphics such as image maps is inherently better than text-based menus or command-level typing (e.g., typing in parameters into a search engine's forms box). Although their comments predate the Web and its omnipresent search engines, Waterworth & Chignell (1989) cited research showing that browsing was best done with a spatial map, while information searching was best done with a text-based index. Symbolic and spatial codes can both be used together. Just as redundant color or shape coding is useful in spatial search on a single page, discrete graphics, as a redundant element to text, aid the learning of a database structure (Merwin, Dyre, & Humphrey, 1995).

SUMMARY

This chapter provides a brief overview of perceptual and cognitive issues in a human factors approach to evaluating and designing Web-based information. Careful consideration of basic features of a Web page, such as color, location, texture, and shape can have a marked effect on the salience of graphics used on a page. Appropriate conspicuousness of graphic features is necessary if these elements are to be effectively organized into an overall design that conveys critical information contained within the page. Beyond taking advantage of perceptual cues, designers can also leverage cognitive factors to help provide recognizable landmarks necessary for effective navigation of a site. Examples of basic design and navigation elements have been outlined, along with emerging areas such as scientific visualization and virtual reality, to demonstrate the necessity of considering perceptual and cognitive elements along with task goals in the design of a Web site.

REFERENCES

Adelson, S. J., Allen, J., Badre, A. N., & Hodges, L. F. (1993). Performance comparison of multiple image depth and shape cues. *International Journal of Human Computer Interaction, 5*, 347–360.

Alty, J. L. (1991). Multimedia—What is it and how do we exploit it? In B. C. Society (Ed.), *People and computers VI* (pp. 31–44). Cambridge, England: Cambridge University Press.

Apple Computer. (1996). Project X Interface. http://www.atg.apple.com/go/projectx/ (accessed October 1996).

Baecker, R. M. (1987). Cognition and information processing theories. In R. M. Baecker & W. A. S. Buxton (Eds.), *Readings in human-computer interaction: A multidisciplinary approach* (pp. 299–307). Los Altos, CA: Morgan Kaufmann.

Baker, P. M., & Wickens, C. D. (1995). *Human factors in virtual environments for the visual analysis of scientific data* (Technical Report No. ARL-95-8/PNL-95-2). Aviation Research Laboratory, Institute of Aviation, University of Illinois at Urbana-Champaign.

Berners-Lee, T. (1995). Style guide for online hypertext. http://www.w3.org/pub /WWW/Provider/Style/Overview.html (accessed October 1996).

Bertoline, G. R., Wiebe, E. N., Miller, C., & Nasman, L. O. (1995). *Engineering graphics communications*. (1st ed.). Chicago: Irwin.

Biederman, I., & Ju, G. (1988). Surface versus edge-based determinants of visual recognition. *Cognitive Psychology, 20*, 38–64.

Brown, T., J. (1991). Visual display highlighting and information extraction. In HFES (Ed.), *Proceedings of the Human Factors Society 35th Annual Meeting* (pp. 1427–1431). Santa Monica, CA: Human Factors and Ergonomics Society.

Carroll, J. M., & Olson, J. R. (1990). Mental models in human-computer interaction. In M. Hellander (Ed.), *Handbook of human-computer interaction* (pp. 45–65). Amsterdam: North-Holland.

Cherry, J. M., Turner, J. M., & Rockwell, G. M. (1992). Designing for diversity: The user interface for a hypermedia information system on a university campus. *Behaviour and Information Technology, 11*, 1–12.

Christ, R. E. (1975). Review and analysis of color coding research for visual displays. *Human Factors, 17*, 542–570.

Cleveland, W. S. (1985). *The elements of graphing data*. Monterey, CA: Wadsworth.

Compuserve, I. (1990). Graphics Interchange Format—Version 89a. http://133.11.25.11/ local/other-faq/gif89a-e.html (accessed October, 1996).

Domel, P. (1995). WebMap: A graphical hypertext navigation tool. *Computer Networks and ISDN Systems, 28*, 85–97.

Dumais, S. T., & Jones, W. P. (1985). A comparison of symbolic and spatial filing. In SIGCHI (Ed.), *Proceedings of CHI '85* (pp. 127–130). New York: ACM.

Ellis, S. R. (1995). Virtual environments and environmental instruments. In K. Carr & R. England (Eds.), *Simulated and virtual realities: Elements of perception* (pp. 11–52). London: Taylor & Francis.

Frey, P. R., Rouse, W. B., & Garris, R. D. (1992). Big graphics and little screens: Designing graphical displays for maintenance tasks. *IEEE Transactions on Systems, Man, and Cybernetics, 22*, 10–20.

Garner, W. R. (1981). The analysis of unanalyzed perceptions. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 119–139). Hillsdale, NJ: Lawrence Erlbaum Associates.

Gillan, D. J., & Lewis, R. (1994). A componential model of human interaction with graphs: 1. Linear regression modeling. *Human Factors, 36*, 419–440.

Haber, R. N., & Wilkenson, L. (1982). Perceptual components of computer displays. *IEEE Computer Graphics and Applications, 2*(3), 23–35.

Hochberg, J. (1981). Levels of perceptual organization. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization: Volume 2. Visual perception.* (pp. 255–278). Hillsdale, NJ: Lawrence Erlbaum Associates.

Hodes, C. L. (1994). Processing visual information: Implications of the dual code theory. *Journal of Instructional Psychology, 21*, 36–43.

Hughes, P. K.,& Creed, D. J. (1994). Eye movement behavior viewing colour-coded and monochrome avionic displays. *Ergonomics, 37*, 1871–1884.

Jubis, R. M. T. (1990). The effectiveness of color-coding in visual displays: Does practice make a difference? In HFES (Ed.), *Proceedings of the Human Factors Society 34th Annual Meeting* (pp. 1519-1523). Santa Monica, CA: Human Factors and Ergonomics Society.

Julesz, B. (1981). Textons, the elements of texture perception, and their interactions. *Nature,* 290, 91-97.

Kaiser, M. K. & Proffitt, D. R. (1989). Perceptual issues in scientific visualization. In International Society for Optical Engineering (Ed.), *Three-dimensional visualization and display technologies* (Vol. 1083, pp. 205-211). Bellingham, WA: International Society for Optical Engineering

Kosslyn, S. M. (1989). Understanding charts and graphs. *Applied Cognitive Psychology, 3*, 185–226.

Levine, R. (1995). Guide to web style. http://www.sun.com/styleguide/ (accessed October 1996).

Lin, R. (1995). An application of the semantic differential to icon design. In G. Perlman, G. K. Green, & M. S. Wogalter (Eds.), *Human factors perspectives on human-computer interaction* (pp. 276–280). Santa Monica, CA: HFES.

Lynch, P. J. (1996). Web style manual. http://info.med.yale.edu/caim/StyleManual_Top.html (accessed October 1996).

Merwin, D. H., Dyre, B. P., & Humphrey, D. G. (1995). Impact of icons and visual effects on learning computer databases. In G. Perlman, G. K. Green, & M. S. Wogalter (Eds.), *Human factors perspectives on human-computer interaction* (pp. 215–219). Santa Monica, CA: HFES.

Meyer, G. W., & Greenberg, D. P. (1980). Perceptual color spaces for computer graphics. *(ACM) Computer Graphics, 14*(3), 254–261.

Morse, A. (1979). Some principles for the effective display of data. *(ACM) Computer Graphics, 13*(2), 94–101.

Padmos, P., & Milders, M. V. (1992). Quality criteria for simulator images: A literature review. *Human Factors, 34*, 727–748.

Petersik, J. T. (1980). The effects of spatial and temporal factors on the perception of stroboscopic rotation simulations. *Perception, 9*, 271–283.

Rao, A. R., & Lohse, G. L. (1993). Towards a texture naming system: Identifying relevant dimensions of texture. In G. M Nielson & D. Bergeron, (Ed.). *Proceedings of visualization '93* (pp. 220–226). New York: IEEE.

Ryan, T., & Schwartz, C. (1956). Speed of perception as a function of mode of representation. *American Journal of Psychology, 69*, 60–69.

Sanford, J., Barfield, W., & Foley, J. (1987). Empirical studies of interactive computer graphics: Perceptual and cognitive issues. In HFES (Ed.), *Proceedings of the Human Factors Society 31st Annual Meeting* (pp. 519–523). Santa Monica, CA: Human Factors and Ergonomics Society.

Silverstein, L. D. (1987). Human factors for color display systems: Concepts, methods, and research. In J. H. Durrett (Ed.), *Color and the computer* (pp. 27–61). New York: Academic Press.

Slack, J., & Conati, C. (1992). Effective graphics: Accessing spatial relations. In British Computer Society (Ed.), *People and computers VII: HCI '92* (pp. 175–191). New York: Cambridge University Press.

Sobol, M. G., & Klein, G. (1989). New graphics as computerized displays for human information processing. *IEEE Transactions on Systems, Man, and Cybernetics, 19*, 893–898.

Swierenga, S. J., Boff, K. R., & Donovon, R. S. (1991). Effectiveness of coding schemes in rapid communication displays. In HFES (Ed.), *Proceedings of the Human Factors Society 35th Annual Meeting* (pp. 1522–1525). Santa Monica, CA: Human Factors and Ergonomics Society.

Thorell, L. G., & Smith, W. J. (1990). *Using computer color effectively*. Englewood Cliffs, NJ: Prentice-Hall.

Treisman, A. M., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology, 12*, 97–136.

Truckenbrod, J. R. (1981). Effective use of color in computer graphics. *(ACM) Computer Graphics, 15*(3), 83–90.

Tufte, E. R. (1983). *The visual display of quantitative information*. Cheshire, CT: Graphics Press.

Valdez, F., Chignell, M., & Glenn, B. (1995). Browsing models for hypermedia databases. In G. Perlman, G. K. Green, & M. S. Wogalter (Eds.), *Human factors perspectives on human-computer interaction* (pp. 147–151). Santa Monica, CA: HFES.

Verplank, W. L. (1990). Graphic challenges in designing object-oriented user interfaces. In M. Hellander (Ed.), *Handbook of human-computer interaction* (pp. 365–376). Amsterdam: North-Holland.

VRML Consortium. (1996). VRML homepage. http://vag.vrml.org/ (accessed December 1996).

Waterworth, J. A., & Chignell, M. H. (1989). A manifesto for hypermedia usability research. *Hypermedia,* 1, 205–231.

Wickens, C. D., & Andre, A. D. (1990). Proximity compatibility and information display: Effects of color, space, and objectness on information integration. *Human Factors, 32*, 61–77.

Wickens, C. D., & Carswell, C. M. (1995). The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human Factors, 37*, 473–494.

Wiebe, E. N. (1994). Evaluation of alternative methods of representing three-dimensional objects on computer displays. In HFES (Ed.), *Proceedings of the Human Factors Society 38th Annual Meeting* (pp. 1326–1330). Santa Monica, CA: Human Factors and Ergonomics Society.

Wiebe, E. N., & Converse, S. A. (1996, October). *Recognition of shape and metric changes in 3–D computer models*. Paper presented at the Human Factors and Ergonomics Society 40th Annual Meeting, Philadelphia, PA.

Zusne, L. (1970). *Visual perception of form*. New York: Academic Press.

Zwaga, H. J. G., & Duijnhower, F. (1995). Influence of a fixed background in a VDU display on the efficiency of color and shape coding. In G. Perlman, G. K. Green, & M. S. Wogalter (Eds.), *Human factors perspectives on human-computer interaction* (pp. 19–23). Santa Monica, CA: HFES.

18 Collaborative Interfaces for the Web

Saul Greenberg
University of Calgary

The 1960s heralded a great leap forward in the way people used computers. During that decade, people moved away from batch processing via punched cards and paper tapes to interactive online dialogues. Yet computers were expensive, and the only way to make interactive systems affordable was through *time-sharing*, a method that allowed many people to interact simultaneously with a single computer. A time-sharing system went to great lengths to give each user the illusion that he or she was the only person working on it (Fig. 18.1a). Personal accounts and access control lists kept people's individual filing systems separate from each other, all in the name of security. Unfortunately, these measures became barriers to collaboration. Information sharing, although possible, was heavyweight. Permissions had to be set; files had to be transferred explicitly from one user's space to another; information had to be included in mail.

The World Wide Web (WWW)—another great leap forward—has changed this process. Simply by making personal or corporate information accessible to a site's Web server, people around the world can distribute and share information with each other. The WWW and its hypermedia structure give the illusion of one large filing system. However, significant barriers to effective collaboration remain. Although information is a shared resource, standard Web browsers are still single-user tools that partition one person from another and offer little support for people to contact each other and engage in conversation over this information (Fig. 18.1b).

In this chapter, I consider how this final barrier to collaboration can be removed from the WWW. I will show how Web browsers can be turned into groupware interfaces that allow people to contact each other, discuss documents, and create artifacts through their displays in real time (Fig. 18.1c). Such systems are now being developed, and there is already a proliferation of Internet-based groupware systems that bring together tools including telephony, address books, text-based chat tools, electronic whiteboards, application sharing, and presentation packages. Commercial examples are the Netscape Conference tools for peer-to-peer communications (http://www.netscape.com/comprod/products/communicator), Microsoft's NetMeeting for multi-point conferencing (http://www.microsoft.com/netmeeting/), and Intel's ProShare which includes video conferencing (http://www.intel.com/proshare/conferencing/). While most commercial systems are not yet well-integrated with browsers, we should expect

this to happen shortly. The recent availability of Javaand powerful plug-in interfaces have made such systems reasonable to implement as part of a Web browser.

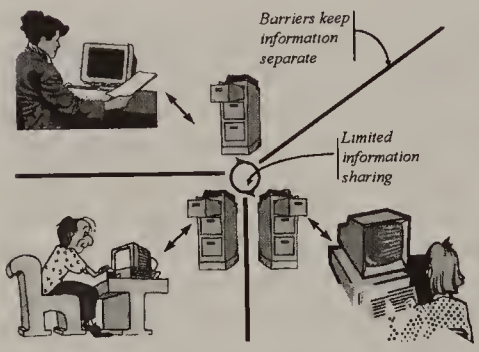


FIG. 18.1a. Traditional time-shared computer systems: Barriers keep people and their information separate from each other.

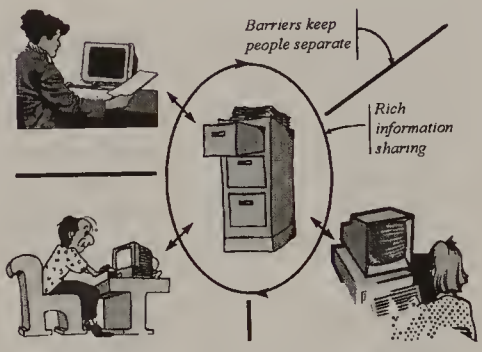


FIG. 18.1b. The WWW: Information is easily shared, but barriers keep people from collaborating on it.

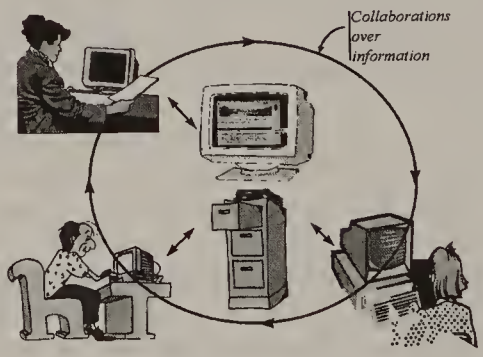


FIG. 18.1c. The WWW and Groupware: People can collaborate on shared information in real time.

FIG. 18.1. From time-sharing to the WWW to Groupware.

Groupware is not a new idea. Engelbart and English (1968), for example, demonstrated voice and video conferencing as well as screen sharing in the late 1960s. Now, an active, decade-old discipline called Computer Supported Cooperative Work (CSCW) studies how people work together (both with and without groupware), and how computers and related technologies can be designed to support effective collaborations. There are specialized conferences on this topic (ACM Proceedings, 1986 – present), edited collections (e.g., Baecker, 1993; Greif, 1988; Greenberg, 1991; Greenberg, Hayne, & Rada, 1995), and even dedicated journals (Schmidt, 1991 – present).

Because the scope of CSCW and its literature is too large to summarize in a few pages, the following section provides an introduction and overview of only one area of groupware and its human factors: shared visual workspaces. The subsequent section provides a case study of a shared visual workspace called GroupWeb, a prototype groupware web browser (Greenberg & Roseman, 1996a, 1996b).

SHARED VISUAL WORKSPACES

In the everyday world, shared physical workspaces (such as whiteboards, control panels, and tabletops) and the artifacts they contain (sketches, controls, documents, structured drawings) act as a stage and offer props for rich person-to-person interaction. Not only are the information artifacts in that space important, but also the ways that interaction over the information is facilitated. For example, people are aware of what others are doing as they are doing it. They can see where others are looking. They can gesture over the surface, can focus other peoples' attention to a part of a workspace; and can guide the presentation of information. They can work individually or collaboratively—even simultaneously—if they want to. Physical workspaces have very specific properties and constraints that afford these things. Creating electronic shared workspaces requires a deep understanding of the affordances of physical workspaces and the way people work over them. This section describes some human factors of a shared workspace.

Studies of Shared Visual Workspaces

Before considering how people interact over a workspace, we first need to know the minimal requirements for effective person-to-person communication. Many studies—for instance, Chapanis, Ochsman, Parrish and Weeks (1972) and Bly (1988)—compared the ways people communicate in real time through technology and contrasted communication conditions such as typed notes, telephone-quality voice, high-fidelity voice, and video links. In all cases, voice communication was found to be important. Although it seems self-evident that people working together in real time over an interactive workspace should have at least a voice channel, many of today's systems provide only a text chat facility, which is inadequate.

A fundamental property of most shared workspaces is the ability to make marks in it. Bly (1988) studied two designers communicating over a drawing surface. From her observations, she asserted that the drawing process—the actions, uses, and interactions on the drawing surface—are as important to the effectiveness of the collaboration as the final artifact produced. She also noticed that allowing designers to share drawing-space activities increases their attention and involvement in the design tasks. When interaction over the drawing surface was reduced, the quality of the collaboration decreased.

Tang (1991) refined Bly's findings through his ethnographic study of small teams solving problems over large sheets of paper. Some of his observations included the following:

- *Orientation.* When people sat around the table, drawings made on the paper were oriented in different directions. Although people had greater difficulty drawing and perceiving the images, orientation provided a resource for facilitating the meeting. Because drawings faced a particular person, a context and an audience were established. Participants' marks that were aligned to an image conveyed support and focus. People working on their own images used orientation as a "privacy" boundary until they were ready to call in the group's attention. The seam when moving between individual and group work was small. (Ishii, Kobayashi, & Grudin, 1993, further developed the idea of "seamlessness".)
- *Proximity.* When participants were huddled close to each other around the table, the drawing played a key role in mediating the conversation.
- *Simultaneous access.* With good proximity, a high percentage (45% to 68%) of people's activity around a work surface involved simultaneous access to the space by more than one person.

Tang then built a descriptive framework to help organize the study of work surface activity, where every user activity was categorized according to what action and function it accomplished. The actions accomplished by user activity included:

- *Listing*, which produces alphanumeric notes that are spatially independent of the drawing,
- *Drawing*, which produces graphical objects, typically a sketch with textual annotations attached to it,
- *Gesturing*, which is a purposeful body movement that communicates specific information, such as pointing to an existing drawing.

The functions accomplished by user activity included:

- *Storing information*, which preserves group information in some form for later recall,
- *Expressing ideas,* which involves interactively creating representations of ideas in some tangible form, usually to encourage a group response,

- *Mediating interaction*, which facilitates group collaboration and includes turn taking and focusing attention.

Tang found that the conventional view of work surface activity—storing information by listing and drawing—constituted only around 25% of all functions. Expressing ideas and mediating interaction comprised the additional 50% and 25% respectively. Gesturing, which is often overlooked as a work surface activity, played a prominent role, about 35% of all people's actions. For example, participants enacted ideas by using gestures to express them, and gestures were used to signal turn taking and focus the attention of the group.

Another important aspect of collaborative work over a shared workspace is gaze awareness, which is a person's ability to monitor the direction of a partner's gaze (Ishii & Kobayashi, 1992). With gaze awareness, one person can see another person's focus of attention. This includes whom the person looks at, whether they make eye contact, whether one person is attending to what the other is doing, and whether people are tracking each other's actions. Gaze awareness also indicates what objects the person attends to in the shared space, and whether both people are looking at the same thing. Through a series of experiments, Ishii, Kobayashi, and Grudin (1992) noticed that gaze awareness helped people facilitate conversation and gesturing, and helped them move between their interpersonal space (where they were looking at each other) and the shared workspace (where they were looking at objects).

Several researchers continued to investigate how people maintain a sense of awareness about who else is in the workspace, where people are operating, and what they are doing (Dourish & Belloti, 1992; Gutwin, Greenberg, & Roseman, 1996). In a physical workspace, people use peripheral vision, auditory cues, and quick glances to track what goes on around them. These kinds of things help people work together more effectively. Gutwin, Greenberg, and Roseman (1996) identified the notion of *workspace awareness*: the collection of up-to-the-moment knowledge a person holds about the state of another's interaction with a workspace. They studied people working together to compose a newspaper layout and noticed the following activities, all afforded by workspace awareness.

- *Mixed focus collaboration*: People shifted their focus back and forth between individual and shared activity, and between different parts of the layout.
- *Lightweight information gathering*: One person gained awareness through rapid glances of another person's working area.
- *Integration of information with previous knowledge.* It was easy to assimilate what a person saw now with what was seen before.
- *Anticipation of another's actions.* A person was aware of what another was going to do next.
- *Using awareness of activity.* A person changed what he or she was doing based on another's activity.

- *Interpreting references.* Sounds, comments, and deictic references of one person were interpreted from another's understanding of where the first was located and what he or she was doing.

Workspace awareness brings another dimension to understanding collaborative interactions. It helps people move between individual and shared activities, provides a context in which to interpret other's utterances, allows anticipation of others' actions, and reduces the effort needed to coordinate tasks and resources. But the notion of awareness is a slippery concept. Gutwin, Greenberg, and Roseman (1996) tried to pin it down by describing it as a list of information elements that people may keep track of when they work with others in a shared space (Table 18.1, column 1). These elements can then be related to questions that people ask themselves during group work (column 2).

Electronic virtual workspaces must emulate the affordances of physical workspaces, if they are to support a group's natural ways of working together. Unlike their physical counterparts, all affordances must be programmed in. This is no easy task. For example, current technology means that some things ordinarily taken for granted are hard to do electronically, such as knowing where others are looking, relating body gestures to items in a workspace, glancing around for awareness, and so on. Consequently, designers must consider and test alternative ways to support what people require; they must understand how the constraints and limitations of technology can be mitigated.

GroupWeb, A Case Study

A prototype groupware Web browser we have developed called GroupWeb (Greenberg & Roseman, 1996a, 1996b) illustrates how knowledge of

TABLE 18.1
Elements of Workspace Awareness

Element	*Relevant Questions*
Identity	Who is participating in the activity?
Location	Where are they?
Activity Level	Are they active in the workspace?
	How fast are they working?
Actions	What are they doing?
	What are their current activities and tasks?
Intentions	What are they going to do?
	Where are they going to be?
Changes	What changes are they making?
	Where are changes being made?
Objects	What objects are they using?
Extents	What can they see?
Abilities	What can they do?
Sphere of influence	Where can they have effects?
Expectations	What do they need me to do next?

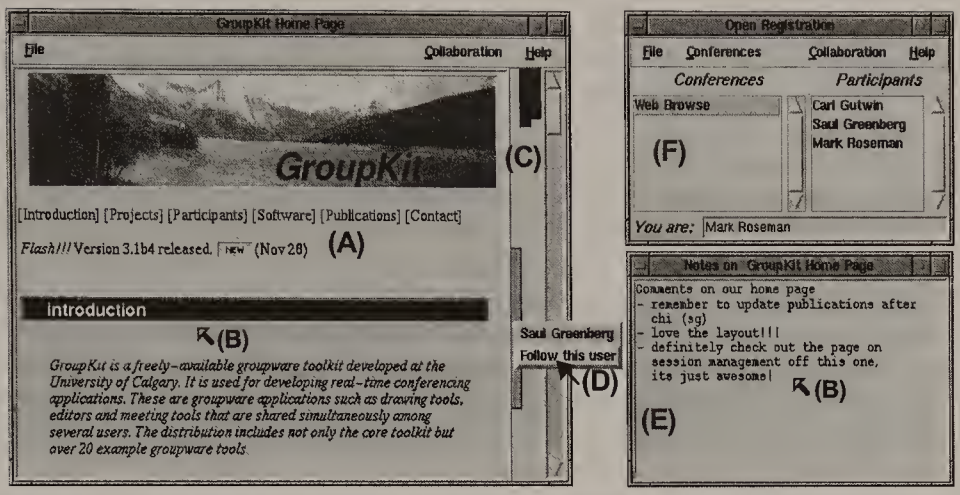


FIG. 18.2. GroupWeb, showing (a) a shared web page, (b) telepointers, (c) multi-user scrollbars, (d) view-slaving controls, (e) a group annotation tool, and (f) a session manager for joining the GroupWeb conference.

workspace activities can be applied to the WWW. A first version is illustrated in Figs. 18.2 and 18.3, and its design rationale follows. Later, we present a redesigned (but unimplemented) second version that addresses some usability problems we had noticed (Fig. 18.4).

Like normal Web browsers, GroupWeb fetches and displays Hypertext Markup Language (HTML) pages. However, several geographically separated people can enter a GroupWeb session via a session manager (Fig. 18.2f). Each participant runs his or her own GroupWeb replica, and the main browser window becomes a shared visual workspace (Fig. 18.2a). GroupWeb currently has the same limitations as normal Web browsers: Thee pages can be viewed but not altered. Nevertheless, GroupWeb can be used as a tool for remote presentations and for discussing documents over distance.

A basic act of a person using a workspace is selecting material (a drawing or page) and bringing it to the group's attention. In GroupWeb, the material is an HTML page that a group member selects by navigating a link. Whenever any person selects a link to a new page, GroupWeb guarantees that all the browsers in the session navigate to that page automatically by instructing all its replicas to fetch the new page, specified by the Hypertext Transfer Protocol (HTTP) address. GroupWeb does *not* enforce turn taking or any other protocol for selecting links. Rather, it relies on people's natural abilities to mediate interaction (as noticed by Tang, 1991) as well as on the affordances of the workspace (such as seeing other people's actions and their gestures) to make sure that an act is reasonable. Although relying on social instead of technical protocols means that conflicting actions can happen, the gain is the ability of the group to follow its own desires and working styles.

How should the page be presented? The simplest thing would be to keep visuals identical across all displays. However, display sizes and personal desires differ, and it may be onerous to impose this constraint onto distributed meeting participants. GroupWeb "relaxes" strict view sharing by permitting windows to be different sizes, and by reformatting text to fit the display nicely (Fig. 18.3). Relaxed view sharing means that all people may not see exactly the same thing (as when orientation differs around a table), but it does provide more flexibility for the way each person wishes to view a page.

A similar question is how scrolling is handled when a page does not fit completely within a window. People sometimes want to pursue individual as well as collective work in a workspace (Gutwin, Greenberg, & Roseman 1996), and they want to be able to go to their own semiprivate space (Tang 1991). GroupWeb allows both independent and synchronized scrolling. With independent scrolling, people can have viewports of different parts of a page, with awareness of others' locations supplied by multi-user scrollbars (Fig. 18.2c). Each colored bar represents how much a particular participant can see, as well as the overlap between views, if any. Synchronized scrolling, on the other hand, automatically aligns viewports. To enable this, a menu on each bar is raised (Fig. 18.2d), with the menu (and the bar's color) identifying the participant it represents. Selecting 'Follow this user' causes the local display to scroll in synchrony with the other.

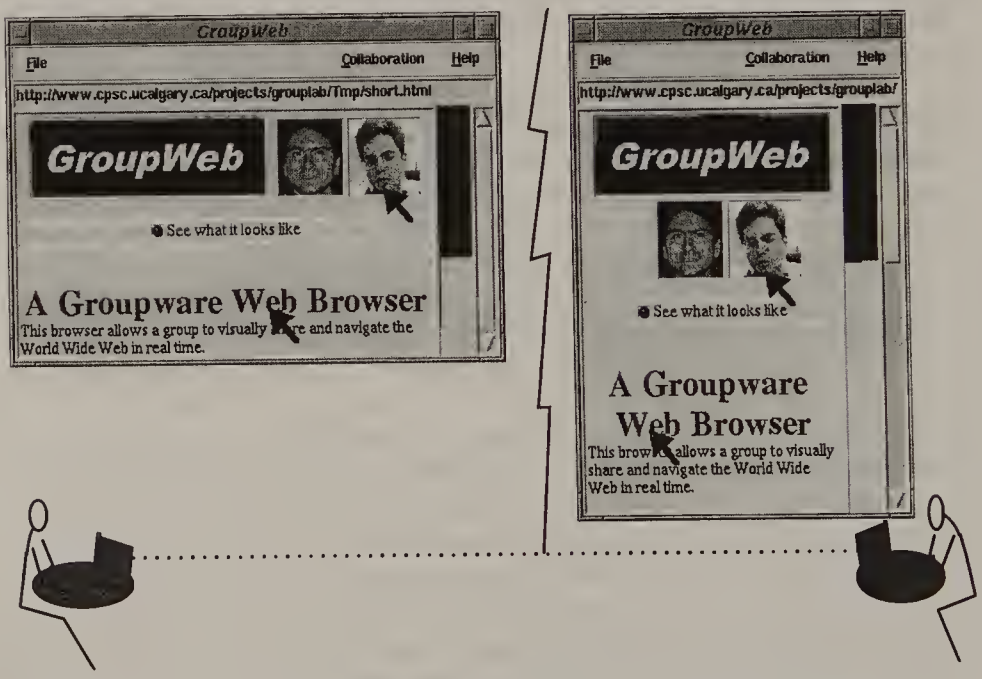


FIG. 18.3. Two participant's view of a GroupWeb session, with contents reformatted to fit the different window sizes. Telepointers are on top of the correct text and images, as they are tied to object coordinates, not Cartesian coordinates.

Because synchronization is one way (unless another person also slaves the view), a local user can still scroll quickly to other parts of the page for quick glances.

Gesturing is critical in a shared workspace. As with many groupware workspaces, GroupWeb uses telepointers as a way to transmit and display gestures (Fig. 18.2b). All participants have their own telepointers, and they use them to enact ideas, to signal turn taking, to focus the group's attention, and to reference objects on the work surface. Telepointers also afford many awareness elements listed in Table 18.1. For example, their presence and motion indicate those who are present, where they are working, what actions they are doing, and what objects they are manipulating. As people often look at their own telepointers as they are using them, telepointers also offer others a limited form of gaze awareness. Yet implementing telepointers can be problematic when the view of the page can be formatted differently on people's displays. In GroupWeb, we solved this problem by attaching telepointers to letter positions rather than to Cartesian coordinates. Thus the pointer is always over the same text on all displays, as illustrated in Fig. 18.3, and most deictic references are correct.

Although the original Web document is not editable, people can attach shared annotations (which Tang, 1991, refered to as a listing action) to any page. GroupWeb includes an annotation tool (Fig. 18.2e), a multi-user text editor implemented as an add-on workspace. Users can simultaneously enter and edit text at any time, which is displayed on all screens. Characters appear on all displays as they are typed, a process useful for expressing ideas. The annotation is automatically keyed to the current WWW page. Changing to a new page clears the editor; returning to an annotated page restores the text annotation in the editor. Annotations can be used for almost anything: group note-taking, collecting comments, suggesting revisions, and chatting. Telepointers work over this area as well.

GroupWeb still has a long way to go, and using it in practice showed some flaws. In spite of being designed around certain human factors requirements, the way some factors are supported is questionable, and other factors are not catered to at all. We mention this because we constantly have to remind ourselves that translating human factors knowledge into systems is not straightforward—usability testing and iterative development *must* be considered as an integral part of the human factors of groupware design.

1. Although GroupWeb users are expected to be in voice contact, they have to establish this connection as a separate act (e.g., conventional or Internet telephony). Because connecting takes time, it would likely deter short collaborative sessions.
2. A multi-user scrollbar was provided to support awareness of the presence of other people and their location in the document. When we tested the scrollbar in a usability study (Gutwin, Roseman, & Greenberg, 1996), we found that people had problems with these scrollbars simply because it was hard to determine exactly where in the

document the other person was, and exactly what portion of the display was visible to both participants. Although the bars could be used to align views, the information required for fine-grained coordination and awareness was too hard to tease out. Other devices that have passed the usability test could be offered instead. One example is a radar overview that presents a miniature of the entire document and draws rectangles and telepointers on the miniature that correspond to people's individual viewports and telepointers. Optionally, it may just be simpler to enforce strict view sharing by eliminating asynchronous scrolling and by guaranteeing same window sizes.

3. Although text listing is possible through the groupware annotation window, People cannot annotate the page directly by drawing graphics and by attaching text marks to objects. Yet this annotation can form 65% of actions in conventional workspaces (Tang, 1991). GroupWeb should include an annotation layer that sits on top of the shared view. Following techniques applied successfully to groupware sketch pads, people should be able to sketch and type simultaneously on top of a displayed page, should see others' drawing actions as they occur, and should allow annotations to be storable.

Figure 18.4 suggests a simpler (but not yet implemented) redesign of GroupWeb. Although much more constrained than the first version, it should be a reasonable tool that allows people to do remote presentations and discussions over documents. In this version, all windows are kept the same size on all screens, and the pages are formatted identically across all displays. Synchronous scrolling eliminates the need for multi-user scroll bars. Telepointers are still present. An annotation layer has been added so that people can both list and draw on top of the image; this layer eliminates the need for an add-on groupware annotation window.

SUMMARY

This chapter has provided a brief overview of the human factors of shared visual workspaces, along with an example of how shared workspaces can be applied in a simple groupware Web browser. The chapter also includes references for those who require a deeper background in CSCW.

Of course, collaborative interfaces in the WWW go far beyond the shared visual workspaces mentioned here. For example, the chapter does not mention the human factors involved in actually getting a groupware session going. People must be aware of who is around and available for conversation (Kraut, Egido, & Galegher, 1988), must initiate their groupware session, and must bring together their shared material. If any of this is hard to do, then it is unlikely that people will bother to start the session unless there is an overwhelming need (Cockburn & Greenberg, 1993). The design of metaphors that bring people and their groupware

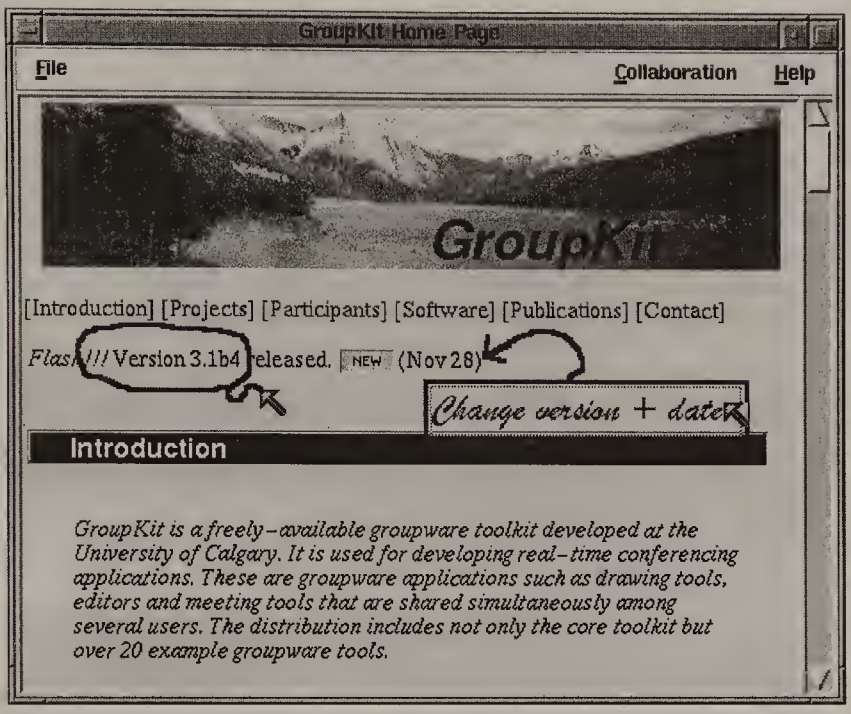


FIG. 18.4. A possible simpler but more constrained redesign of GroupWeb. All windows are kept the same size on all screens. Synchronous scrolling eliminates the need for multi-user scroll bars. An added annotation layer eliminates the need for the add-on groupware editor.

together is just as critical as the design of the groupware application. Other issues include security (so that sessions remain private), access control (so that a group can decide who can enter a groupware session), the melding of individual and group work (so that people's individual work can be brought in and out of groupware sessions), the melding of asynchronous and synchronous work, and so on. As well, collaborative interfaces should be designed to support specific tasks. For example, the demands of group authoring have a different (but perhaps overlapping) set of human factors requirements than those found in shared visual workspaces (Rada, 1996).

The spreading of the Internet, the popularity of the WWW, the ubiquity of WWW browsers, and the recent availability of Java mean that collaborative interfaces will be built for the WWW. Fortunately, these collaborative interfaces do not have to be crude first efforts with poor usability. We can bring to their design the rich human factors literature and experiences in CSCW and groupware.

REFERENCES

Baecker, R. (1993). *Readings in groupware and computer-supported cooperative work: Assisting human-human collaboration*. San Mateo, CA: Morgan-Kaufmann.

Bly, S. (1988). A use of drawing surfaces in different collaborative settings. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (pp. 250–256), Portland, OR: ACM.

Chapanis, A., Ochsman, R., Parrish, R., & Weeks, G. (1972). Studies in interactive communication: I. The effect of four communication modes on the behaviour of teams during cooperative problem solving. *Human Factors, 14*(6), 487–509.

Cockburn, A., & Greenberg, S. (1993). Making contact: Getting the group communicating with groupware. In *Proceedings of the ACM Conference on Organizational Computing Systems* (pp. 31–41). Milpitas, CA: ACM.

Dourish, P., & Bellotti, V. (1992). Awareness and coordination in shared workspaces. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (pp. 107–114), Toronto, Ontario: ACM.

Engelbart, D., & English, W. (1968). A research center for augmenting human intellect. In *Proceedings of the Fall Joint Computing Conference* (Vol. 33, pp. 395–410). Montvale, NY: AFIPS Press.

Greenberg, S. (Ed.). (1991). *Computer supported cooperative work and groupware*. London: Academic Press.

Greenberg, S., Hayne, S., & Rada, R., (Eds.). (1995). *Groupware for real-time drawing: A designer's guide*. Berkshire, England: McGraw-Hill.

Greenberg, S., & Roseman, M. (1996a). GroupWeb: A groupware web browser. In *Video Proceedings of ACM Conference on Computer-Supported Cooperative Work*, Boston: ACM.

Greenberg, S., & Roseman, M. (1996b). GroupWeb: A WWW browser as real time groupware. In *Companion Proceedings of the ACM Conference on Human Factors in Computing Systems* (pp. 271–272). Vancouver, BC: ACM.

Greif, I. (Ed.). (1988). *Computer-supported cooperative work: A book of readings*. San Mateo, CA: Morgan-Kaufmann.

Gutwin, C., Greenberg, S., & Roseman, M. (1996). Workspace awareness in real-time distributed groupware: Framework, widgets, and evaluation. In R. Sasse, A. Cunningham, & R. Winder (Eds.) *People and computers XI Proceedings of HCI'96 BCS Conference on Human Computer Interaction*, (pp. 281–298). Springer-Verlag.

Gutwin, C., Roseman, M., & Greenberg, S. (1996). A usability study of awareness widgets in a shared workspace groupware system. In *Proceedings of ACM Conference on Computer Supported Cooperative Work*, Boston: ACM.

Ishii, H., & Kobayashi, M. (1992). ClearBoard: A seamless medium for shared drawing and conversation with eye contact. In *ACM Conference on Human Factors in Computing Systems* (pp. 525–532). Monterey, CA: ACM.

Ishii H., Kobayashi M., & Grudin, J. (1993, October). Integration of interpersonal space and shared workspace: Clearboard design and experiments. *ACM Transactions on Information Systems*.

Kraut, R., Egido, C., & Galegher, J. (1988). Patterns of contact and communication in scientific collaboration. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work* (pp. 1–12). Portland, OR: ACM.

Proceedings of the Conference on Computer Supported Cooperative Work. Biannually since 1986, ACM.

Rada, R. (Ed.). (1996). *Groupware and authoring.* London, England: Academic Press.

Schmidt, K. (Ed.). (1991). *Computer Supported Cooperative Work: The Journal of Collaborative Computing.* Kluwer.

Tang J. (1991, February). Findings from observational studies of collaborative work. *International Journal of Man Machine Studies*, *34*(2), 143–160.

19 A Zooming Web Browser

Benjamin B. Bederson
James D. Hollan
Jason Stewart
David Rogers
David Vick
University of New Mexico

Laura Ring
Eric Grose
Chris Forsythe
Sandia National Laboratories

The World Wide Web (WWW) is becoming increasingly important for business, education, and entertainment. Popular Web browsers make access to Internet information resources relatively easy for users. Simply by clicking on a link, a new page of information replaces the current one on the screen. Unfortunately, after following several links, people can have difficulty remembering where they have been and which navigating links they have followed. As people's collections of Web pages grow and as more information of interest populates the WWW, effective navigation becomes an issue of fundamental importance.

People with diverse talents, interests, and experiences are online via the WWW. Many are novices with little experience using computers. A simple click of the mouse can bring a user from a friend's home page to unknown destinations across the world. This process can be disorienting becausc, traditionally, following a cross-reference meant shuffling across the library to find another volume. Although time consuming, this movement reinforced the transition that was taking place. The difficulties that novice users confront can be instructional to developers. Experts may not have as much difficulty, but they experience similar cognitive burdens and may simply have a higher threshold before they experience similar difficulties.

Although the immediacy of traversing information links offers many advantages, it can also make it difficult for users to maintain an intuitive sense of where they are, and how they got there—a classic problem of hypertext systems. Part of the problem can be attributed to windows-based interfaces. Current window systems do not readily support showing more than a few pages at a time. In addition, each page is usually in a separate window with no depiction of relations to other windows. Popular Web browsers, like other applications built according to windowing philosophies, also have this problem,

although they do offer some methods to aid navigation by keeping track of interesting sites—usually in hierarchical sets of hotlists or bookmarks.

Many human factors issues must be considered (Kellogg & Richards, 1995). Some of these have been addressed by proposed alternatives and extensions to browsers. WebMap is a browser extension that shows a graphical relation between Web pages (Doemel, 1994). Each page is represented by a small circle that can be selected to display the actual page. The links between pages are colored to indicate information about links, such as whether the link is to a different server or whether the destination page has already been read. These graphs may be saved and used by others. WebBook is a three-dimensional Web browser from Xerox Palo Alto Research Center (Xerox PARC); it allows multiple pages to be viewed simultaneously and supports collections of pages in books that replace bookmark files (Card, Robertson, & York, 1996).

Web pages often have links to pages that link back to the starting point. As it is difficult to visualize these cycles, many Web visualizations are based on hierarchies extracted from a graph of the WWW. Some interesting work has focused on alternative visualizations (Mukherjea, Foley, & Hudson, 1995). Furnas and Zacks (1994) showed that multitrees can be used to represent a collection of hierarchies sharing parts of the underlying data. One application of multitrees is visualization of bookmarks from multiple individuals.

Another approach to visualizing large information spaces applicable to Web browsing and navigation involves techniques to show detail at particular areas while maintaining context. One general approach, fisheye views (Furnas, 1986), has been extended with graphics (Sarkar & Brown, 1994), three dimensions (MacKinlay, Robertson, & Card, 1991), hyperbolic representations (Lamping, Rao, & Pirolli, 1995), animation (Chang & Ungar, 1993), and zooming (Bederson & Hollan, 1994; Bederson et al., 1996; Perlin & Fox, 1993). Other techniques include exploiting a large virtual space (Donelson, 1978), using lenses or filters (Bederson et al., 1996; Stone, Fishkin, & Bier, 1994), and visualizing two-dimensional layouts (Bederson & Hollan, 1994; Bederson et al., 1996).

One approach described by Furnas (1995) provides a basic framework for discussing what he termed "view navigable" structures. He distinguished "view traversal," the simple mechanical process of moving around interactively by clicking on things in each view, from "view navigation." View navigation adds to view traversal the requirements of effectively finding paths to targets and discerning the right things to click on to get to desired places. A major interface challenge is how to give users some indications of what is on the other side of a link so that they can more effectively judge which path to follow.

The Montage system (Helfman, 1997) exemplifies one promising approach to this problem. Montage is a new way to surf the WWW by watching hundreds of images from Web pages. For example, Montage can create a dynamic image map, a cycling montage of images from a traversal of a Web site, to harvest the images from pages at that site. Selecting an interesting image from this montage drives a regular Web browser to the associated page. Montage is a fast way to visualize the landscape of the WWW because it utilizes the human capacities to recognize relevant images and to deal with many images at once. The Montage approach is interesting because a large percentage of Web images

have content that is either explicitly textual or in some way relevant to the text on the associated Web page. Montage changes the nature of Web navigation. By removing images from their associated Web pages, Montage flattens the hypertext structure of the Web.

Finally, navigating the WWW in a zooming environment, can show Web pages in context. In the following sections, we describe our zooming Web browser, discuss user studies comparing it to Netscape Navigator from Netscape Communications Corporation (http://www.netscape.com), and show our improvements based on an initial study; the result was a substantial speedup in the Pad++ Web browser.[2]

A Zooming Web Browser

Navigating the WWW is a struggle between focus and context. As users browse or search the WWW, the need for detailed views of specific items conflicts with the need to maintain a global view of context and history of traversal. The global view is important because it is often necessary to access previously visited pages, either as a starting point to find new information or for comparison.

We are building a new zoomable Web browser using Pad++, a substrate for building multiscale dynamic user interfaces (UIs) (Bederson & Hollan, 1994; Bederson et al., 1996). Pad++ provides an extensive graphical workspace where dynamic objects can be placed at any position and at any scale. Pad++ supports panning and zooming. Zooming can involve simple geometric scaling, which we term *context-sensitive rendering*, in which rendering of objects can vary based on factors in addition to scale, such as context of the task or complexity of the information being displayed. Pad++ is built as a component for Tcl/Tk, a scripting language and UI library (Ousterhout, 1994).

Pad++ allows Web pages to remain visible at varying scales so viewers may examine many pages at once. In addition, Pad++ allows users to zoom in and out and thus offers explicit control of how much context is viewed at any time. To orient themselves, users can simply zoom out to view a number of Web pages. To get more detailed views of a particular page, they can zoom in. We think that this variable-scale contextual display of Web pages can provide important support for navigation. When using our Web browser, users navigate a space filled with familiar objects, not iconic representations of these objects.

Interaction

Users interact with the Pad++ Web browser in much the same manner as traditional browsers. Clicking on a link loads in a new page.

The Forward and Back buttons from traditional Web browsers are slightly modified to allow access to all the pages previously visited, not just those that were visited recently. Traditionally, a Back button goes to the page from which the current page was accessed. We replaced the Back button with Previous and Parent buttons. The

---

[2] Pad++ and the zooming Web browser are available to the public. To find current information, look at <http://www.cs.unm.edu/pad++>, or send email to pad-info@cs.unm.edu

Previous button does the same thing as the original Back button, but the Parent button goes to the page that originally linked to the current page.

There is also a button called "Show All" which zooms out just enough so that all pages are visible on the screen. Then, a single click on any page focuses on the page and brings it to a readable size.

The left mouse button is used for following hyperlinks by clicking on them and for panning by dragging on the surface. Pressing on the middle button zooms in until the button is released. Pressing on the right button zooms out until the button is released. Pad++ zooms around the mouse cursor. In this way, users can control which point the view is zoomed around by moving the mouse while zooming.

Our Web browser combines a basic focus-driven layout with automatic zooming and panning to support navigation. The software allows users to select a focus page. This selection animates the page to occupy a larger section of the display. Pages farther from the focus page get increasingly smaller, so that a graphical fisheye view results. Figures 19.1 and 19.2 show snapshots of the Pad++ Web browser with different pages focused on the area where the view is zoomed out so that many pages can be seen. Figure 19.3 shows a zoomed-in view where a single page may be read. Ancestor and descendant pages are visible on the sides; they can be clicked on to bring them into focus.

Clicking on a link brings up a new page, adds it to the tree of pages, and causes the tree to restructure itself. Unlike other Web browsers that immediately replace the current page with a new page, the restructuring process is animated so that users can understand how the tree is being updated. The animation helps maintain object constancy, and the graphical depiction of links highlights relations between pages. The new page becomes the current focus and is moved to the center of the screen at a size suited for viewing. Users may designate any existing page to be the current focus by clicking on it.

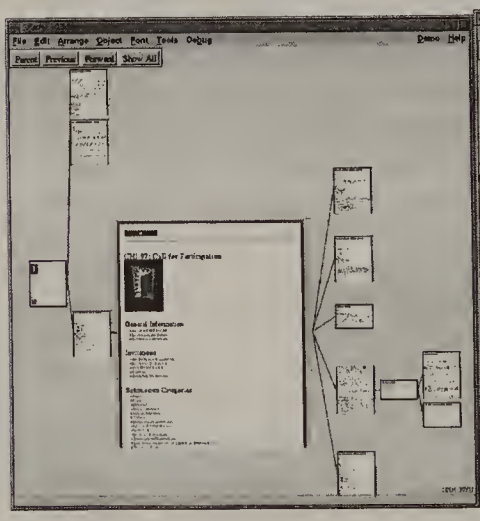


FIG. 19.1. A snapshot of the Pad++ Web browser.

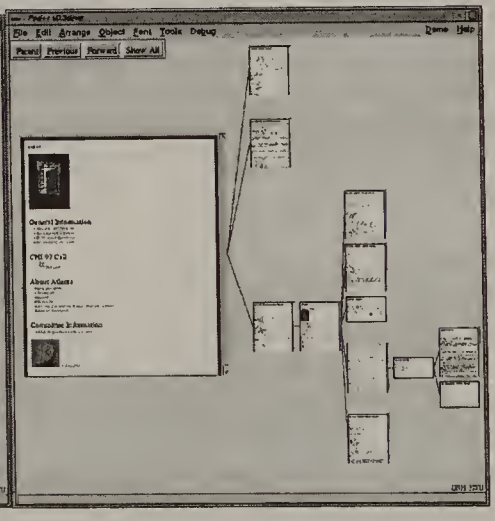


FIG. 19.2. The same Web pages with a different focus.

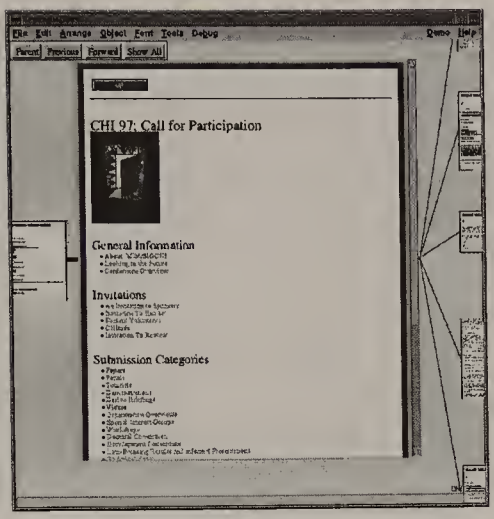


FIG. 19.3 Zoomed into a single page.

Clicking on a link to a page that has already been loaded focuses on that page. Even if originally accessed from another part of the tree, the page's position in the tree is not changed.

As in earlier fisheye displays, our basic layout function assigns a degree-of-interest value to each page in the tree on the basis of a page's distance from the focus page. We define the distance to be the shortest path between two pages. This value is then used to determine the size of each page. Because long pages have such a big effect on the layout, we added scrollbars to pages longer than a given size. This way, users can still see the structure of the pages, but long pages are not as disruptive.

Usability Testing

We have hypothesized that the Pad++ Web browser improves users' abilities to retrieve, comprehend, and integrate data from the WWW, intranets, or Web-based databases. We think that the ability to see multiple pages simultaneously, and to learn the structure of the WWW facilitates many tasks. We chose to study the Pad++ Web browser in a focused technical domain that allowed measurement of times to answer simple questions. In addition, this domain is representative of the use of a Web browser in a corporate environment.

THE STIMULI

We designed a hypothetical database containing the evolution of a manufacturing process over a series of six custom product builds, hereafter referred to as *versions*. The entire stimulus set consisted of 31 Web pages (see Fig. 19.4). An Overall page contained information on the database and links to the Design and Manufacturing Process pages and links to the Budget and Schedule pages for each of the six versions. Six Budget and

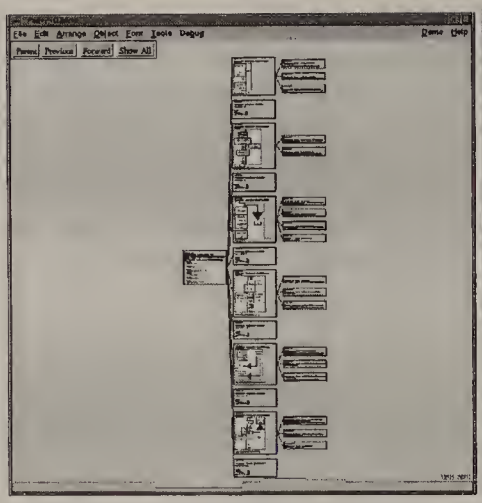


FIG. 19.4. Overview of all stimuli.

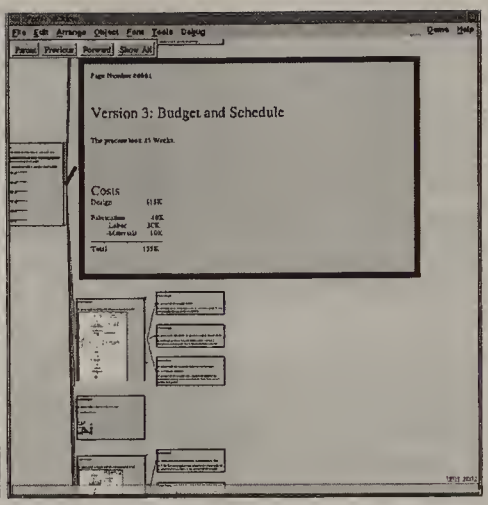


FIG. 19.5. A *budget and schedule* page from stimuli.

Schedule pages (see Fig. 19.5) gave an overall estimate of time and budget for each version. Six Design and Manufacturing Process pages contained the process flowcharts for each version (see Fig. 19.6). The process flowcharts were image maps with links to notes about the process. The final 18 pages contained these notes.

Task 1 required subjects to navigate to specified pages in the database and enter a page number found on the top of each page. For example, the subjects were asked to enter the page number of Version 5's Budget and Schedule page. This task tested speed of navigation in the browser.

Tasks 2 and 4 tested users' ability to retrieve and comprehend information about the whole database. In Task 2, subjects were asked to count the number of versions with certain features or where certain events occurred. For example, they were asked to indicate the number of versions that had a new step added. Task 4 required users to identify the version in which a certain feature or event occurred. For example, subjects were asked which version had a problem with fabrication.

Task 3 tested users' ability to retrieve and integrate information found in different pages of the database. Subjects were asked to compare different versions. For example, they were asked to indicate which version, out of versions one, two, and three, had the highest material cost. This task required subjects to look at the Budget and Schedule page for each of the three versions and compare their material costs.

Experimental Conditions

Four experimental conditions were examined: Pad++ Prebuilt, Pad++ Interactive, Netscape Regular and Netscape Graphic. In the Pad++ Prebuilt condition, the entire network of pages was loaded into Pad++ before beginning the task. In the Pad++ Interactive condition, only the Overall page was loaded and the network of pages was built as users answered the questions. In the Netscape Graphic Condition, an image map of the entire database was created from a screen capture in Pad++. This condition was

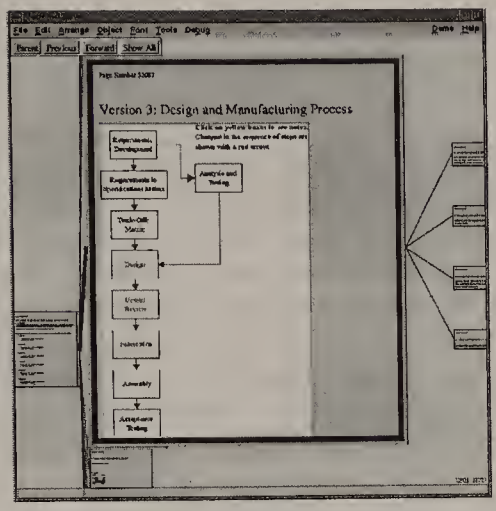


FIG. 19.6. A *design and manufacturing* process page.

added to see whether Pad++'s navigational tools caused improvement in performance independent of that caused by the visual representation of the database. The image map used contained a snapshot of the pages depicted in Fig. 19.4.

Training

If a subject reported using Netscape less than 2 hours a week, they were asked to complete a brief training task. The task, taking about 10 minutes, consisted of browsing through a set series of pages from UNM's Web site. During this task, subjects were instructed on the use of the Forward and Back buttons and on how to click on hypertext and image maps to follow links.

All but one subject had never used Pad++, although a few had seen brief demonstrations of the software. Three training tasks were used to familiarize the subjects with panning and zooming in Pad++ and using the Pad++ browser. Training Tasks 1 and 2 were used to familiarize subjects with a three-button mouse. In Training Task 1, subjects were asked to zoom in and out. In Training Task 2, they were asked to zoom, pan, and select objects. The subjects performed both tasks until they were able to perform 10 operations without error. The third training task asked subjects to open a series of pages in UNM's Web site and then to navigate to different open pages. During the training, subjects were informed of different strategies to use when navigating in the Pad++ environment.

Subjects

The experiment was carried out at UNM. There were 32 participants, but only 30 participants yielded usable data. The 2 participants whose data were not used had over 50% of the questions incorrect. All the subjects had completed at least 2 years of college, and all used a computer at least 2 hours a week. Most subjects (73%) used a computer more than 1 hour a day. Almost all the subjects (93%) had browsed the WWW before,

and 80% of subjects browsed the WWW at least 1 hour a week. Most (90%) were students, and the rest (10%) were professionals.

Procedures

Before beginning the experimental tasks, the subjects were trained if needed and then asked to read through an instruction page. The instruction page explained the task and database and had links to example pages. Subjects were instructed to click on all links and look at all the examples. They then carried out trials for the four experimental tasks. In one window, the browser was open with the appropriate initial pages, and in another window, the questions were displayed. A program written in Tcl displayed the questions and recorded the time spent on each question and the subjects' responses. After each question, the entire screen went blank so that users could not see the browser and database pages. The subjects were instructed to click on a bar to proceed. This action started the timer and displayed the question and the browser window. Subjects clicked on another bar to enter their responses and to stop the timer. After the experimental task, subjects filled out a questionnaire with questions about the browser, their computer usage, and their demographic information.

Results

The percentage of subjects who answered each question correctly was determined, and if over 20% missed the question, it was considered too difficult and discarded for further analysis. Out of 24 questions, 3 were thrown out for this reason. Of these questions, 2 were in Task 2 and 1 was in Task 3.

The dependent variable was the average time it took subjects to answer each question. Only the times for questions answered correctly were used in calculating the average time. An average time to correct response was calculated for each of the four tasks. The average times collapsing across tasks were also calculated. The number of correct responses ranged from 16 out of 21 to 21 out of 21, with most subjects having 1 or fewer incorrect responses

We performed a one-way analysis of variance on the average time to correct response for each task. We had limited statistical power because only 30 subjects yielded usable data. Because of this lack of power, we decided to use an alpha level of 0.10 instead of 0.05 for each test.

Figure 19.7 shows a graphical representation of the average time to correct response for each task and collapsed across tasks. The overall F test was only significant for Tasks 1 and 2, with an alpha level of $p < 0.10$, $F_{3,26} = 2.74$, $p = 0.064$ and $F_{3,25} = 2.66$, $p = 0.070$, respectively. These results were significant because subjects in the Pad++ Interactive Condition took longer to carry out the tasks than did the subjects in the Netscape Graphic condition.

These results were largely attributable to Pad++'s technical problems that made it run considerably slower than Netscape. On the basis of the qualitative data collected by observation and from subjects' responses to the questionnaire, we made improvements to the Pad++ browser as described in the next section

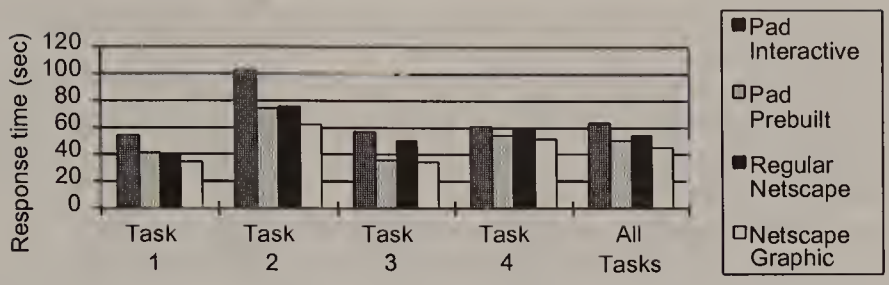


FIG. 19.7. Average time to correct response by task and browser condition.

The speed at which Pad++ was running caused some discontent, but many comments suggested that subjects felt the interface to be intuitive and fun to use. Some quotes indicating positive aspects of Pad++ included:

ZOOM is fun!

The fact that one can get the "whole picture" in varying degrees of closeness is fantastic.

I love the zoom in and out. I also liked Show All very much.

The layout of the pages made the task very easy once one figured out that the pages were laid out in the order listed on the parent page. With the color scheme used, the zooming browser allowed answering of some questions merely by panning and examining the pages from a low resolution view.

Easy to find what I need to know. I did not have to search very much.

I liked the tree structure, but disliked the occasional delays in updating the display. Once I knew where the pages were in the tree structure, I could quickly find what I wanted.

Another interesting observation made by the experimenter during testing was that subjects varied in their use of panning and zooming available in Pad++. Because of their familiarity with browsing in Netscape, they seemed reluctant to try some of the other features available to them in Pad++. They stuck to what they knew, despite the fact that in training, the experimenter had encouraged them to use all available navigational tools. Furthermore, there seemed to be a positive relation between subjects' use of these navigation features and task performance. In future studies, it would be interesting to quantify this difference between subjects and to see whether it is positively correlated with task performance.

IMPROVEMENTS MADE TO THE PAD++ BROWSER

Improvements made to the Pad++ browser after the first experiment included:

- Making Pad++ run roughly twice as fast by using an 8-bit instead of a 16-bit X-server.
- Changing a feature so that the page would be centered when it was in focus. When a page was initially brought into focus it was aligned with the left side of the screen. Some subjects had pointed out that this alignment left them feeling disoriented because of a loss of context.
- Improving fonts from mediocre to high-quality system fonts.
- Adding a feature so that users needed only to click again on the page to bring it out of focus and make it small again. During testing, subjects had pointed out that they wanted to be able to defocus a page as well as focus it..
- Making pages thinner so that more context could be seen, a change made possible by the improved font quality.
- Fixing a cursor bug that occurred only when we tested the software with novice users. This bug had caused the cursor to freeze on a zooming cursor if a user pressed two mouse buttons at once.
- Adding a feature in Pad++, so that the hand cursor changed into an arrow over a link, when some subjects had indicated that they missed the Netscape feature of the cursor's changing when it was over a link.

EXPERIMENT 2

Following the first experiment, we realized that we needed to make some changes to evaluate Pad++ more fairly. First, we made the just-mentioned improvements. Second, as almost all the subjects in the Netscape conditions were regular Netscape users, we decided that to evaluate Pad++ in a comparable manner, we should use subjects who were familiar with Pad++.

We ran the Pad++ Prebuilt condition again with 7 Pad++ users and the same procedures as discussed previously. All 7 subjects had worked as part of the Pad++ development team; all were familiar with navigation in the Pad++ environment and did not need to be trained.

Results

A one-way analysis of variance was performed to compare the new subjects' performance to those subjects in the Netscape conditions of the previously described study. The average times to correct response for each task and averaged across all the tasks are displayed in Fig. 19.8.

The overall F-test approached significance for Task 1, and for the average across all tasks, $F_{2,19} = 2.33$, $p = 0.124$ and $F_{2,19} = 2.43$, $p = 0.115$, respectively. Furthermore, tests of the apriori hypothesis that Pad++ would do better on average than the Netscape conditions yielded a $p = 0.0599$, $F_{1,19} = 4.00$ for Task 1 and a $p = 0.088$, $F_{1,19} = 3.23$ for the average time to correct response across all tasks.

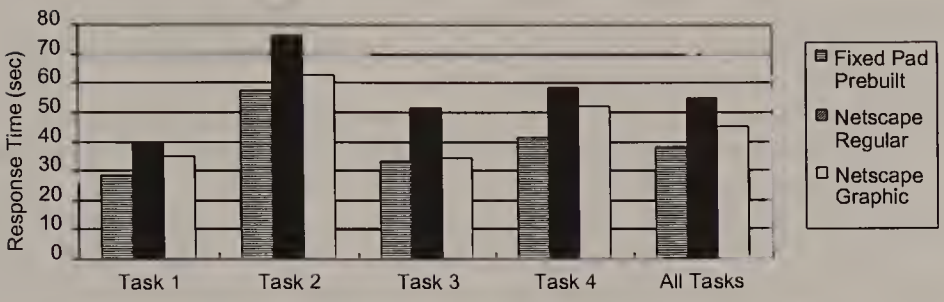


FIG. 19.8. Average time to correct response for improved Pad++ versus the Netscape conditions by task.

SUMMARY

These results have demonstrated that subjects can perform this browsing task 23% faster on average using the Pad++ browser than using Netscape. Many subjects also found the Pad++ browser intuitive and fun to use. In a corporate setting, use of the Pad++ browser could increase workers' productivity by decreasing the time it takes them to access historical information from various corporate databases. Furthermore, as using the Pad++ browser allows users to access such information in a fun and intuitive manner, it could lead to an increased use of historical information with consequent improvements in product and process quality.

Because the WWW has become an important and widely used resource, it is crucially important to address its usability. We have shown one promising technique based on zooming that supports improved Web navigation.

ACKNOWLEDGMENTS

We acknowledge generous support from the Defense Advanced Research Project Agency's (DARPA's) Human-Computer Interaction Initiative (Contract #N66001-94-C-6039). This work was also supported by the U.S. Department of Energy under Contract DE-AC04-96AL85000.

We appreciate the work of our collaborators at the New York University Media Research Lab, especially Jon Meyer, in helping us develop Pad++.

REFERENCES

Bederson, B. B., & Hollan, J. D. (1994). Pad++: A zooming graphical interface for exploring alternate interface physics. In *Proceedings of User Interface Software and Technology (UIST '94;* pp. 17–26). New York: ACM.

Bederson, B. B., Hollan, J. D., Perlin, K., Meyer, J., Bacon, D., & Furnas, G. (1996). Pad++: A zoomable graphical sketchpad for exploring alternate interface physics. *Journal of Visual Languages and Computing, 7*, 3–31.

Card, S. K., Robertson, G. G., & York, W. (1996). The WebBook and the Web Forager: An information workspace for the World-Wide Web. In *Proceedings of Human Factors in Computing Systems (CHI '96;* pp. 111–117). New York: ACM.

Chang, B., & Ungar, D. (1993). Animation: From cartoons to the user interface. In *Proceedings of User Interface Software and Technology (UIST '93*; pp. 45–55). New York: ACM.

Doemel, P. (1994). WebMap—A graphical hypertext navigation tool. In *Second International Conference on the World-Wide Web* (pp. 785–789). Chicago: International WWW Conference, http://141.142.3.76/SDG/IT94/Proceedings/ Searching/doemel/www-fall94.html.

Donelson, W. C. (1978). Spatial management of information. In *Proceedings of Computer Graphics (SIGGRAPH' 78*; pp. 203–209). New York: ACM.

Furnas, G. W. (1986). Generalized fisheye views. In *Proceedings of Human Factors in Computing Systems (CHI '86*; pp.16–23). New York: ACM.

Furnas, G. W. (1995). *Effectively view-navigable structures*. Paper presented at the Human Computer Interaction Consortium Workshop (HCIC '95), Snow Mountain Ranch, Colorado.

Furnas, G. W., & Zacks, J. (1994). Multitrees: Enriching and reusing hierarchical structure. In *Proceedings of Human Factors in Computing Systems (CHI '86*; pp.330–336). New York: ACM.

Helfman, J. I. (1997). Montage: Expanding Web Accessibility Through Image Surfing. Submitted to *Proceedings of User Interface Software and Technology (UIST '97)*. New York: ACM.

Kellogg, W. A., & Richards, J. T. (1995). The Human Factors of information on the Internet. In J. Nielsen (Ed.), *Advances in human computer interaction* (pp.1–36). Norwood, NJ: Ablex.

Lamping, J., Rao, R., & Pirolli, P. (1995). A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings of Human Factors in Computing Systems (CHI '95*; pp. 401–408). New York: ACM.

Mackinlay, J. D., Robertson, G. G., & Card, S. K. (1991). The perspective wall: Detail and context smoothly integrated. In *Proceedings of Human Factors in Computing Systems (CHI '91*; pp. 173–179). New York: ACM.

Mukherjea, S., Foley, J. D., & Hudson, S. (1995). Visualizing complex hypermedia networks through multiple hierarchical views. In *Proceedings of Human Factors in Computing Systems (CHI '95*; pp. 331–337). New York: ACM.

Ousterhout, J. K. (1994). *Tcl and the Tk toolkit*. Reading, MA: Addison-Wesley.

Perlin, K., & Fox, D. (1993). Pad: An alternative approach to the computer interface. In *Proceedings of Computer Graphics (SIGGRAPH' 93*; pp. 57–64). New York: ACM.

Sarkar, M., & Brown, M.H. (1994). Graphical fisheye views. *Communications of the ACM, 37*(12), 73–84.

Stone, M. C., Fishkin, K., & Bier, E. A. (1994). The movable filter as a user interface tool. In *Proceedings of Human Factors in Computing Systems (CHI '95*; pp. 306–312, New York: ACM.

Author Index

A

Abelson, R. P., 48, 49
Abrams, I., 36
Adams, A. J., 29
Adelson, S. J., 158, 233
Albrecht, J.E., 49, 51
Allen, J., 233
Alty, J. L., 227
Ambrose, H. W., 19
Anderson, J.R., 54
Andre, A. D., 231
Anstis, S. M., 28, 29, 38
Apple Computer, 157, 164, 234
Archer, J. A., 17, 18
Askari, N., 35

B

Bacon, D., 256
Badre, A. N., 233
Baecker, R. M., 227, 243
Baggett W. B., 53, 54
Baker, P. M., 226, 232, 233
Balota, D.A., 45, 48
Barfield, W., 234
Barnard, P., 208
Barnes, V., 200
Barnett, S. A., 17, 18
Bassok, M., 54
Beam, P., 201
Becklen, R., 39
Bederson, B.B., 256, 257
Bellotti, V., 245
Benjamin, B., 138
Bennett, E. L., 21
Berlyne, D. E., 17
Berners-Lee, T., 225
Bernstein, S., 210, 220
Bertoline, G. R., 231
Bertus, E.L., 54
Besner, D., 36
Biederman, I., 33, 35, 234
Bier, E.A., 256
Birke, L. I. A., 17, 18
Blake, R., 29
Bloom, C. P., 49, 51
Blundon, W., 155
Bly, S., 243, 244
Bobo, G., 20
Boff, K. R., 229
Boice, R., 20
Boltman, A., 80
Borges, J. A., 141, 144, 148, 149
Bower, G.H., 53, 54
Boyes-Braem, P., 34
Bransford, J., 48, 49
Britton, B. K., 52
Brown, M. H., 256
Brown, T., J., 229, 231
Burden P., 100
Burkhard, B., 18

C

Card, S.K., 256
Carpenter, P. A., 48
Carr, T. H., 33
Carroll, J. M., 196, 217, 227
Carswell, C. M., 231
Cato, S., 205
Champoux, M., 21
Chang, B., 256
Chang, F.R., 53
Chapanis, A., 243
Chase, P., 34
Chen, C., 36
Cherry, J. M., 228
Chi, M.T.H., 54
Chignell, M., 234, 235, 236
Christ, R. E., 226, 231
Chrysler, S. T., 47, 49
Clark, H.H., 53, 55
Clark, L. F., 51
Cleveland, W. S., 230
Cockburn, A. J. G., 166, 250
Comber, T., 138
Compuserve, Inc., 229
Conati, C., 225
Connell, 89
Converse, S. A., 232
Corbett, A.T., 53, 54
Cowan, P. E., 18
Crane, R. S., 55
Creed, D. J., 231

Cross, D. R., 36
Crown, J., 201
Csikszentmihalyi, M., 196
Cuqlock-Knopp 67

D

Dahlgren, K., 49
Dansereau, D. F., 36
Davies J., 100
Deaton, J.A., 49
December, J., 115, 130
DeJong, G., 49
Dell, G. S., 53
DeMello, C., 138
Deni, R., 20
Dern, D. P., 98
Dijkstra, K., 55
Dillion, A., 200, 202
Dirik, A., 176
Dix , A. J., 207
Doemel,, P., 256
Domel, P., 234
Dominguez, C., 64
Domjan, M., 17, 18
Donelson, W. C., 256
Donovon, R. S., 229
Dosher, B. A., 54
Dougherty, T., 35
Dourish, P., 245
Druin, A. M., 75, 78, 80
Duffy, S. A., 52
Duijnhower, F., 230, 231
Dumais, S. T., 158, 236
Duncan, J., 38
Durup, M., 19
Dyer, M. G., 49
Dyre, B. P., 235, 236

E

Egido, C., 250
Eisner, E. W., 76
Elkind, J., 90
Ellis, S. R., 233
Ellison, L., 176
Engelbart, D., 243
England, D., 208
English, W., 243
Erwin, H., 20
Eterovic, V. A., 21

F

Faletti, J., 36
Farquhar, A., 158
Ferchmin, P. A, 21
Fincher-Kiefer, R., 55
Fisher, K. M., 36
Fishkin, K., 256
Fisk, A. D., 39
Fletcher, C. R., 47, 49, 51
Flores d'Arcais G.B., 45
Fodor, J. A., 26
Foley, J., 234, 256
Foster, R. W., 3
Fox, D., 256
Frey, P. R., 228, 234
Friedland, L., 176
Friedrich, F. J., 36
Furnas, G. W., 158, 256
Furuta, R., 165

G

Galegher, J., 250
Garner, W. R., 226
Garris, R. D., 228
Gelade, G., 231, 232
Gernsbacher, M. A., 45, 49, 53
Gentner, D., 186
Gibson, J. J., 25, 27, 31
Gillam, B., 26
Gillan, D. J., 230
Givon T., 49
Glaser, R., 54
Glass, A. L., 35
Gleitman, H., 38
Glenberg, A. M., 51
Glenn, B., 235
Glenn, C. G., 46, 235
Gold, C., 36
Gold, R., 196
Golding, J.M., 51, 54
Goldman, S.R., 51
Gomez, L. M., 158
Gomez-Pinnilla, F., 20
Goodwin, L. G., 20
Gould, B., 29
Gould, J. D., 200, 201
Graesser, A. C., 45, 46, 48, 49, 50, 51, 53, 54
Gray, W. D., 34

Gray, P., 208
Green, G. C., 20
Greenberg, D. P., 231
Greenberg, S., 166, 243, 245, 246, 248, 249, 250
Greenough W. T., 21
Greenspan S. L., 54
Greif, I., 243
Grice, H. P., 55
Grinberg, J. G. A., 75
Grischkowsky, N., 200, 201
Gruber, T., 158
Grudin, J., 245
Gulgoz, S., 52
Guthrie, K. M., 20
Gutwin, C., 245, 246, 248, 249
GVU, 153, 156, 171

H

Haake, J. M., 165
Haber, R. N., 226, 227, 233, 234
Hailman, J. P., 27
Halldorson, M., 53
Halliday, M. A. K., 49, 116
Halliday, M. S., 17
Hanneman, J., 165
Hartley, J., 202
Hartson, R., 196
Hasan, R., 49, 116
Haviland, S. E., 53, 55
Hayne, S., 243
Hebert, M., 21
Helfman, J. I., 256
Heller, H., 138, 155, 157
Hempel, M., 21
Henik, A., 36
Hilton D. J., 54
Hix, D., 196
Hochberg, J. E., 32, 227
Hodes, C. L., 227, 234
Hodges, L. F., 233
Hof, R. D., 176
Hoffman, J. E., 33
Hoft, N. L., 1
Hollan, J. D., 256, 257
Holley, C. D., 36
Horton, W., 1,2, 200, 203
House, D., 78, 80
Hudson, S., 256
Hughes, P. K., 231
Human Engineering Guidelines for Mgmt Information Systems, 66
Humphrey, D. G., 235, 236
Humphreys, G. W., 38
Husebye-Hartmann, E., 54

I

Igbaria, M., 196
Ignacio, A., 1
Inglis, I. R., 17, 18
Ishii, H., 245

J

Jagielo, J. A., 17
Johnson, B. K., 54
Johnson, C., 207, 210, 212, 213, 215, 216
Johnson, D. M., 34
Johnson, M., 48, 49
Jones. S. R. A., 166
Jones, W. P., 236
Jonides, J., 29, 38
Ju, G., 234
Jubis, R. M. T., 231
Juell P., 108
Julesz, B., 26, 226
Just, M. A., 48

K

Kaiser, M. K., 233
Kavanagh, J., 210, 212, 215, 216
Kawasaki, G., 168
Kay, A., 75
Kay, R. H., 115
Keefe, D. E., 55
Keenan, J. M., 54
Kehoe C. M., 104, 117
Keller, H., 17
Kellogg, W. A., 256
Kiesler, S., 186
Kimchi, R., 32
Kinchla, R. A., 33
Kintsch, E., 48, 49, 52, 53
Kintsch, W., 49, 50, 52, 53
Klein, G., 228
Klin, C. M., 55
Kobayashi, M., 245
Koenig, O., 26
Kosslyn, S. M., 26, 230

Kraut, R., 12, 250
Krebs, J. R., 18
Krile T., 108
Kreuz, R. J., 52
Kuhmann, W., 208

L

Lambiotte, J. G., 36
Lamping, J., 256
Land, E. H., 30
Landauer, T. K., 158
Langston, W. E., 49, 50, 51
Latremouille, S. A., 201
Lauren, U., 130
Lawicka, W., 20
Lazarro, J., 87, 88
Leon., M., 20
Lepoivre, M., 20
Lepper, M.R., 196
Levine, R., 137, 225
Lew, G. S., 202
Lewis, C. H., 185
Lewis, M.W., 54
Lewis, P., 78
Lewis, R., 230
Lickorish, A., 200, 201
Lin, R., 227
Lipson, J., 36
Loftus, G. R., 35
Lohse, G. L., 226
Long, D. L., 51
Long, M., 77
Lorch, E. P., 52
Lorch, R. F., 49, 51, 52
Lowgren, J., 130
Ludolph, F., 186
Lund, A. M., 129
Lynch, G., 21
Lynch, J. J., 20
Lynch, P. J., 137, 225

M

Mack, R. L., 165
Mackinlay, J. D., 256
Mackworth, N. H., 35
MacLeod, C. M., 36
Magliano, J. P., 49, 50, 53, 54, 55
Malone, T. W., 196
Mandler, J. M., 48
Manning, J., 186
Mannoni B., 106
Markowitz, H., 21
Marr, D., 26, 27, 29
Maunsell, J. H. R., 26
May, J., 102, 208
McAleese, R., 202
McCall, M. A., 20
McCormick, E. J., 29
McGowan, S., 208
McDaniel, M. A., 55
McKoon, G., 50, 53, 55
McNamara, D. S., 53
Menzel, E. W., 19
Mervis, C. B., 34
Merwin, D. H., 235, 236
Metzgar, L. H., 18, 19
Meyer, D. E., 35
Meyer, G. W., 231
Meyer, J., 256
Microsoft NetMeeting, 241
Milano, D., 170
Milders, M. V., 233
Miller, C., 231
Miller, J. S., 17
Millis, K. K., 45, 54
Miller, R., 115
Moore, C., 200
Morales, I., 141, 144, 148, 149
Morimoto, H., 21
Morris, D., 20
Morrow, D. G., 54
Morse, A., 230
Moryadas, A., 36
Mottershead, B. E., 20
Moxey, L. M., 249
Mukherjea, S., 256
Mukhopadhyay, T., 186
Murphy, A. R., 88
Murray, J. D., 55
Muter, P., 201
Myers, J. L., 52, 55

N

Nabkel, J., 115, 116
Nadel, L., 17, 18, 19
Nadeau, D. R., 25
Nasman, L. O., 231
Navon, D., 32, 33
Neely, J. H., 36

Neal, L., 196, 216
Neisser, U., 39
Newsome, W. T., 26
Nielsen, J., 3, 4, 130, 138, 139, 143, 158, 165
Nieto, S. M., 20
Nissen, M. J., 38
Noel, A., 36
Norman, D. A., 116
Norman, J., 33
Norton, P., 75, 76
Nygren, E., 161, 162

O

O'Brien, E. J., 49, 51
Ochsman, R., 243
O'Keefe, J., 17, 18, 19
Ogden, W. C., 38
Olson, J. R., 227
Oostendorp, H., 47
Oracle Corporation, 176
Ousterhout, J. K., 257

P

Padmos, P., 233
Palmer, S. E., 32, 34
Papert, S., 75
Paquette, D., 21
Parrish, R., 243
Patterson, H., 36
Pausch, R., 196
Pejtersen, A.M., 214
Pelli, D. G., 29
Perlin, K., 256
Petersik, J. T., 234
Pezdek, K., 35
Piernot, P., 158
Pinker, S., 33
Pirolli, P., 256
Pitkow J.E., 104, 117
Platt, J. G. A., 75, 80
Pollatsek, A., 33, 36, 37
Posner, M. I., 38
Potter, M., 36
Potts, G. R., 54
Poucet, B., 19
Prescott, J., 21
Proffitt, D. R., 233
Przybyslawske, J., 20
Public Service Commission of Canada, 94, 95

R

Rabinowitz, J. C., 35
Rada, R., 36, 243, 251
Ramachandran, V. S., 31
Randall, N., 115, 130
Raney, G. E., 37
Ranney, M., 54
Rao, A. R., 256
Ratcliff, R., 50, 53, 55
Ratner, J. A., 111
Rayner, K., 36, 37, 45
Reichel, F. D., 31
Reimann, R., 54
Reinhardt, V., 21
Renner, M. J., 17, 18, 19, 21
Revlin, R., 53
Reynolds, K., 35
Reynolds, S. B., 36
Rice, J., 158
Richards, J. T., 256
Riha, J. R., 46
Risden, K., 54
Ritchie, B. G., 55
Rivenburgh, R., 158
Rivers, D., 138, 155, 157
Roberts, R. M., 52
Robertson, G. G., 256
Robertson, S. R., 256
Robson, J. G., 29
Rock, I., 27, 30, 32
Rockwell, G. M., 228
Rodríguez, N. J., 141, 144, 148, 149, 150
Rosch, E., 34
Roseman, M., 243, 245, 246, 248, 249
Rosenzweig, M. R., 21
Rosmaita, G., 92, 94, 95
Rouse, W. B., 228
Rudnicky, A. I., 208, 220
Russel, P. A., 17, 18
Ryan, C., 186
Ryan, T., 234

S

Sanders, M. S., 29
Sanford, A. J., 234
Sarason, B. R., 18
Sarason, G., 18

Sano, D., 138, 155, 158, 160, 165, 170
Sarkar, M., 256
Schafer, A. B., 19
Schank, R. C., 54
Scherlis, W., 186
Schiffman, S. J., 196
Schmidt, K., 243
Schneider, W., 39
Schuh, E., 155
Schumacher, R. M., 202
Schvaneveldt, R. W., 35
Schwartz, C., 234
Secco, S., 51
Sekuler, R., 29
Seltzer, C. P., 17, 19
Seminerio, M., 187
Sengul, C. J., 53
Shafir, E, 115, 116
Shank, R. C., 48, 49, 54
Shiffrin, R. M., 39
Shillito, E. E., 18
Shneiderman, B., 186
Shinjo, M., 52
Silverstein, L. D., 226
Singer, M., 49, 51, 53
Skelly, T., 196
Slack, J., 225
Slater, P., 176
Slugowski, B.R., 54
Smith, M. C., 36
Smith, R. E., 18
Smith, S. W., 158
Smith, W. J., 226, 230, 231
Sobol, M. G., 228
Solomon, C., 75
Songer, N. B., 53
Spear, N. E., 17
Spear, P. D., 20
Sperry, L., 51
Spinelli, J., 21
Spoehr, K. T., 75
Spring, C., 36
Spyridakis, J. H., 202
Stacy, E. W., Jr., 35
Standal, T. C., 202
Stein, N. L., 46
Stone, M.C., 256
Stotts, P. D., 165
Strehlo, K., 167, 168
Suh, S., 51, 54
Swierenga, S. J., 229

T

Tang J. C., 244, 245, 247, 248, 249, 250
Tannenbaum, D., 90
Tauscher, L., 166
Taylor, L., 1
Teal, S.L., 208, 220
Thaggard, P., 54
Theeuwes, J., 38
Thiel, D., 196
Thielcke, G., 20
Thinus-Blanc, C., 19
Thomas, J. C., 196
Thorell, L. G., 226, 230, 231
Thorndyke, P. W., 65
Thornton, R., 36
Thovtrup, H., 130
Thüring, M., 165
Tinker, R. F., 76
Toates, F. M., 17, 18
Todd, J. T., 31
Tolman, E. C., 63
Tomaso, N., 20
Trabasso, T., 48, 49, 51, 54
Treisman, A., 38, 231, 232
Treurniet, W. C., 201
Truckenbrod, J. R., 230
Trueman, M., 202
Tufte, E. R., 231
Tumaso, N., 20
Tulving, E., 36
Turkle, S., 76, 77
Turner, J. M., 228
Tzelgov, J., 36

U

Ungar, D., 256
Uscher, N., 75, 78, 80

V

van den Broek, P., 48, 49, 51, 54
Van Dijk, T.A., 48, 49, 52
Van Essen, D. C., 26
Valdez, F., 235
Vanderheiden, G. C., 87, 92
Vince, M. A., 20
Virtual Architecture Group, 233
Vora, P. R., 158, 165

Varnhagen, C. K., 51
Verplank, W. L., 228, 230
Vetter R., 108
Voss, H. G., 17

W

Wakefield, D., 92
Waller, R., 116
Walters, C., 167
Wang, L. S., 29
Ward, L. M., 33
Waterworth, J. A., 234, 235, 236
Weeks, G., 243
Weinman, L., 160, 163
Wells, J., 21
Whalen, S., 196
Whetstone, T., 35
Whitaker, L. A., 66, 67
Whitney, P., 55
Wichansky, A. M., 176
Wickens, C. D., 63, 226, 231, 232, 233
Wiebe, E. N., 231, 232, 234
Wieckowski, T. J., 196
Wieringa, D., 200
Wilhite, S. C., 202
Wilkins, A. J., 29
Wilkenson, L., 226, 227, 233, 234
Wilson, S., 157, 159, 160, 163
Wolfe, J. M., 33
Wong, L., 29
Wright, P., 200, 201

Y

Yantis, S., 29, 38
York, W., 256

Z

Zacks, J., 256
Zimmer, U., 20
Zusne, L., 227
Zwaan, R. A., 45, 47, 49, 50, 55
Zwaga, H. J. G., 230, 231

Subject Index

A

Accessibility, 87-95
Acquisition
  functional information, 19
Adaptive technology, 87-88, 93
Aerial perspective, 31
Aesthetic appreciation, 52
Age, 29
Americans with Disabilities Act, 88
Animation, 230, 258-259
Annotation, 249
Applet, 183-185
Applications, 126-128, 130
  functionality, 186
  limitations, 186
  web-enable, 175
Archie, 98-99
Argument overlap, 47, 49-50
Assistive technology, 87
Attention, *see also* Selective attention
  content, 192
  dividing, 39
  resources, 48, 230
  sources of information, 192
  visual search, 38-39
Audio, 90, 93
Automation, 108

B

Behavioral homeostasis, 20
Binocular disparity, 27, 31
Blinking, 230
Bookmark, 65, 70
Bottom-up processing, 49
Braille translators, 93
Browser, 102, 121, 166, 187, 202, 208, 228, 243, 246-250, 255-256
  functionality, 7, 180
  interface, 111, 115
  settings, 202
  stream, 233-234
  testing, 7
  text-only, 91
  usability, 112-113
Bulletin board, 104
Buttons, 90, 149, 180, 194, 258

C

Card sorting technique, 158, 190-191
Case study, 246-250
CD ROM, 102, 104
Children, 75-84, 196
Cinematic approach, 216-218
Classification errors, 192
Client models, 176
Cluster studies, 192
Coarsework, 98-101, 105-106
Cognitive
  map, 235
  models, *see also*, Mental models, 113
  psychology, 64
Coherence, 49, 51
Coherent relationships, 49
Cohesion, 49
  cues, 116
Collaborative
  learning environments, 78
  tools, 84
Color 226, 229-230, 231-232
  contrast, 30
  pattern, 229
Common gateway interface (CGI), 108
Common ground, 55-56
Common look and feel, 70, 90, 123, 130, 140, 164
Communication, 77, 82
  person-to-person, 243
Computer-Supported Cooperative Work (CSCW), 10-11, 243
Conceptually-driven processes, 26-28, 33, 35-36
Content analysis, 157-159
Context, *see also* Scene, 202
  sensitive rendering, 257
Contrast, 90, 229
  color, 30
  lightness, 30
  sensitivity function, 29-31
Cooperative communication, 55-56
Cross platform, 176
Crowding effect, 29, 38
Cultures, 77-79
  corporate, 121
Curriculum delivery, 101, 111, 116, 118

D

Database, 64, 70, 187
  administration, 183-185
Data-driven processing, 26-28
Data
  dimensions, 232
  representation, 230-233
Degree-of-interest value, 259
Demand crisis, 207
Demographics, 112
Depth cues, 31
  pictorial, 31
Dialogue, 139
Diffused image, 30
Disequilibrium, 20-21
Document
  distribution, 199
  printed, 225
  structure, 202
Drawing, 84, 244
Dynamic imaging, 233

E

Education, 6-7, 64, 75, 77, 97-108, 111
Element of surprise, 216-217
E-mail, 6, 79-80, 82, 84, 104, 181
Enterprise manager, 183-185
Entertainment, 195-196
Environment
  classroom, 117
  collaborative learning, 77-78
  enrichment, 21
  hyperlinked, 39, 130
  structured, 65-66
  unstructured, 66-68
  virtual, 69
Event-indexing model, 50
Experience, 64, 66, 155, 228
  engaging, 190
Exploration, 5, 17, 25, 63-65, 75
Expositions, 47

F

File size, 201, 210-212
File Transfer Protocol (FTP), 98, 105
Focus-driven layout, 258
Footer, 139-140
Formatting, 202-204
Form perception, 33
Forms, 91, 108
Frame of reference, 64, 66
Frames, 71, 90, 91, 235

G

Gaze awareness, 245
Geographical location, 212-213
Gestalt psychology, 32
Gesturing, 244-245, 249
Given-new strategy, 55
Global
  coherence, 51
  connections, 50
  precedence effect, 32-33
  relationships, 49
  structure, 32
Glyphs, 232
Good continuation, 32
Gopher, 99-100, 105
Graphics, 10, 88, 90-93, 140, 160, 163, 201, 225-236
  border, 229
  images, 194
  primitives, 226-227
Groupware, 241-250
GroupWeb, 246-250
Guidelines, 10, 88-90, 94-95, 117, 121-136, 137-151, 153, 160, 164, 189-195, 196-197, 203, 208-220

H

Hardware, 70-71, 94, 155
Heading, 139, 202-203
Helper application, 208
Hierarchical cluster analysis, 190
History, *see also* World Wide Web
  computing, 241
  sites, 65
Home page, 139
Human Factors
  reviews, 126
Hypertext Markup Language (HTML)
  functionality, 179-180
  limitations, 180-181

I

Icon, 65, 70, 91-92, 114, 140, 194, 235
Illusions, 26-27

Illusory contour, 31
Index, 140, 144, 146-149, 202
Inferences, 53-55
  anaphoric, 53
  associations, 54
  explanatory, 54
  predictive, 54-55
Information
  clusters, 158
  encoding, 227
  quality, 104
  repositories, 228
  sharing, 241
  spaces, 256
  structure, 234-236
Infrastructure, 208-210, 213
Intellectual property rights, 107-108
Interaction, 108, 159, 205, 245
Internet, 77
  delivery, 199
  exploration, 21
  learning, 17-18
  teaching, 98
  traffic, 140, 144, 168
  transatlantic, 103
  usability, 12-13
InterOffice, 177-182
Interposition, 31
Intranet, 121, 155-156
Invariant properties, 25

J

Jargon, 115
Java, 235, 243

L

Landmarks, 63, 65, 70, 116, 193-194, 235
Layout, 66, 161, 201-202, 214-216
Learning, 18, 20-21, 39, 75, 77
  curve, 117-118
  distance, 83-84
  online environments, 78-84
  problem-centered, 75-76, 78, 81
Library, 102, 106
Limited resources model, 227
Linear perspective, 31
Links, 69, 91, 139-140, 229, 235, 256
  naming, 144-146
Load testing, 167
Local
  connections, 50
  relationships, 49

M

Machine intelligence, 64
Maintenance, 168
Management structure, 157
Managing content, 156
Manipulable objects, 228
Market, 87-88
Memory
  long-term, 48, 52
  processes, 19
Mental
  model, 67, 157-158, 228, 234
  representations, 47
  space, 69
Menu
  bar labels, 113
  text-based, 236
Metaphors, 161, 186, 194
Methodology, 154-169
Mirror sites, 212-213
Modem, 210
Montage, 256-257
Motivation, 75
Multimedia, 10, 104, 160, 163
Multiple handicaps, 89

N

Narratives, 47
Navigation, 5-6, 10-12, 19, 63-71, 94, 100, 113, 139-140, 144-149, 180-181, 189-190, 192-195, 201-202, 229, 235-236, 255-257
  bars, 202
  buttons, 149
  cues, 67-68, 70
  keyboard-based, 90, 94
  strategies, 67-68
Newsgroup, 104
Network, 97
  Computer (NC), 182-183, 187
  Computing Architecture (NCA), 176
  local, 114
Novelty, 17-18
Novices, 64, 66, 111-118

O

Object identification, 35
Object-oriented technologies, 156
Object recognition *see* Recognition
Occlusion, 31
Online
  documentation, 116
  environments, 77-78
  satisfaction questionnaire, 166
  services, 155
  technology, 76
Organic approach, 190
Organizational signals, 51
Ownership, 107-108

P

Pad++, 257-265
Paper, 9, 115-116, 199
  conversion, 200-205
  copy, 140, 204
  documents, 199
Pattern recognition, 33-34
Perceptual
  errors, 26-27
  proximity, 231-232
  separation, 232
Platform independence, 121
Pragmatics, 55-56
Priming effect 36
Principles, 69-71
Process analysis, 159
Production quality, 210
Protocols, 97
Psycholinguistics, 45
Public relations, 156

Q

Questionnaire, 166, 196

R

Random-dot stereograms, 26-27
Readability, 200-201
Reading, 36-38
  comprehension, 116
Real time, 77-78, 93
Recall, 227
Recognition
  by components, 33
  object, 33-35
  pattern, 33-34
  scene, 34
  word, 35-36
Redundant information channels, 229
Referential overlap, 50
Remote control, 182
Representation
  mental, 47
  surface, 47
  text, 47
Representational frameworks, 34-35
Response time, 10, 82, 93, 99-103, 153, 156, 167, 186, 201, 207-208, 210-216
  search, 38, 162-163
Route finding, 65
Route knowledge, 63

S

Satisfaction, 216
Scene
  context, 34-35
  recognition, 35
Screen
  magnifier, 90
  reader, 90
  resolution, 201-202
  size, 202
Scrolling, 91, 139, 248
Search
  engines, 64, 69, 104, 106-107, 114
Selective attention, 38-39
Shared workspaces, 243-246, 249
Situation awareness, 64, 69-70, 235
Situation model, 47, 50, 54
Size constancy, 31
SmartCard, 182
Software design
  traditional, 4
Spatial
  cues, 116
  frequency, 29-30
  frequency analysis, 29
  information, 19
  learning, 194
  layout, 19
  map, 236
  sensitivity, 29-30
Spatially continuous, 51

Spontaneous exploration, 17
Stereopsis, 27, 31
Stereotypes, 203
Students
  college, 98, 111, 116
  elementary school, 76, 78-84
Style Guide, 7-9, 122-124, 127, 130, 164
  audience, 123, 130
  authors, 122-123
  noncompliance, 128-129
Subwebs, 158-159
Supply crisis, 207
Survey, 101-108, 122, 125, 153, 155
  knowledge, 63-64

T

Table of contents, 70, 202
Tables, 88, 90, 93
Tags, 91-92
Task-directed
  group, 214
  users, 214-215
Task-motivated users, 215
Team problem-solving, 244
Telepointer, 249
Television, 182, 187
Templates, 159
Text
  alternative, 90, 92, 94
  comprehension, 45, 48, 51-56
  expository, 46
  formatting, 203
  narrative, 46
  procedural, 46
  representations, 47
  text-based displays, 37-38
Texture, 26, 226
  gradient, 31
Three-dimensional
  form, 227
  realism, 234
  shape, 34
Thumb-nail sketches, 215
Title, 70
Top-down processing, 49
Training, 98-99
Transfer of training, 186

U

Universal design, 89
Unix, 104
Usability, 113-114, 116, 138, 153-154, 156
  engineering, 7-8, 137, 154, 165, 189-190, 196
  evaluation, 138
  heuristics, 138-144, 150
  perception, 113
  remote, 166
  researchers, 195
  study, 158, 165-166, 193
  testing, 8-11, 91-92, 118, 138, 140-143, 165-167, 179-183, 218-219, 259-265
Usage
  characteristics, 155
  environments, 155
  scenarios, 177-178, 182
  trends, 168-169
User-centered
  approach, 190, 228
  design, 176-177
  hierarchy, 190-191
User Interface
  design, 129, 178
  issues, 4
  testing, 129
Usenet, 104

V

Video, 9-10, 90, 207-220
Virtual
  reality, 233-234
  workspaces, 246
Visual
  acuity, 28-29
  field, 38
  search, 38-39
Visualization, 84, 234, 255-265
Virtual Reality Markup Language (VRML), 34, 233

W

WebBook, 256
WebMap, 256
WebTV, 170
Web Attitude Measure, 111, 115
Web shows, 194
Window
  management, 179
  systems, 255
Workspace awareness, 245-246

Work surface, 245-246
World Wide Web, 79
content, 186, 195-197, 210-212
design, 39, 46, 52, 55-56, 64-66, 69, 87, 94-95, 137-138, 143, 145, 153-171
development, 8-9, 127
history, 1, 100
instruction, 116
interface, 1
page-creating applications, 91
personalization, 169
publishing, 1-2, 108, 130, 157
research tool, 80
statistics, 3, 90, 153
testing, 88
usability, 2-4, 8

Z

Zooming, 257-265

HUMAN FACTORS AND WEB DEVELOPMENT

Usability engineers will find this book helpful in understanding how their expertise can be applied to Web design and will be inspired to get involved in any Web development in their client or employment organizations. Web designers of all kinds will find this book to be a good introduction to the usability issues involved in Web design, as well as an inspiration to pursue other usability engineering resources to support their efforts. It should not take another 20 years for the Web development community to catch up to the traditional development community in terms of the application of usability engineering principles and methods to product development. This book should contribute to the process.

—Deborah Mayhew

Edited by

Chris Forsythe
Eric Grose
Julie Ratner

ISBN 0-8058-2823-0

90000

9 780805 828238

ISBN 0-8058-2823-0