
Literacy and Computers

The Complications of Teaching and
Learning with Technology

Edited by

Cynthia L. Selfe and Susan Hilligoss

The Modern Language Association of America
New York 1994

© 1994 by The Modern Language Association of America
All rights reserved
Printed in the United States of America

Library of Congress Cataloging-in-Publication Data

Literacy and computers : the complications of teaching and learning
with technology / edited by Cynthia L. Selfe and Susan Hilligoss.
p. cm. — (Research and scholarship in composition)
Includes bibliographical references and index.
ISBN 0-87352-579-5 (cloth) — ISBN 0-87352-580-9 (pbk.)
1. Computers and literacy—United States. 2. Language arts—
United States—Computer-assisted instruction. 3. Reading—United
States—Computer-assisted instruction. 4. English language—
Composition and exercises—Study and teaching—United States—
Computer-assisted instruction. I. Selfe, Cynthia L., 1951–
II. Hilligoss, Susan, 1948– . III. Series.

✓LC149.5.L575 1993 4
371.3'34—dc20 93-26959

Published by The Modern Language Association of America
10 Astor Place, New York, New York 10003-6981

Printed on recycled paper

Contents

Preface to the Series	ix
Introduction <i>Cynthia L. Selfe and Susan Hilligoss</i>	1
PART I: Changing the Contexts of Literacy Instruction	
Overview Reading, Writing, and Thinking in an Age of Electronic Literacy <i>William Costanzo</i>	11
Chapter 1 The Politics of Literacy and Technology in Secondary School Classrooms <i>Paul J. LeBlanc</i>	22
Chapter 2 Blinding Insights: Classification Schemes and Software for Literacy Instruction <i>Gail E. Hawisher</i>	37
Chapter 3 Interpreting the Discourses of Technology <i>Ellen L. Barton</i>	56
Response Literacy, Technology, and Teacher Education <i>Jane Zeni</i>	76
PART II: Extending Literacy through Computer Networks	
Overview Reading and Writing on Computer Networks as Social Construction and Social Interaction <i>Ann Hill Duin and Craig Hansen</i>	89

Chapter 10

The Effect of Hypertext on Processes of Reading and Writing

Davida Charney

Texts and Hypertexts

Most people conceive of *text* as a collection of ideas that a writer has carefully selected, framed, and organized into a coherent sequence or pattern in hopes of influencing a reader's knowledge, attitudes, or actions. A key element in this conception of text, from the perspective of both writers and readers, is structure. Linguists and discourse analysts have identified a host of structural patterns that writers work with (and, more frequently, against) at every level of text production, from small units such as sentences and paragraphs all the way to grand structures that describe entire texts, such as sonnets, fairy tales, résumés, or policy arguments (Halliday and Hasan; van Dijk; Fahnestock and Secor). Indeed, readers depend on such patterns to identify a text's genre, anticipate its development, and integrate its parts. Studies of reading comprehension confirm that readers understand and learn most easily from texts with well-defined structures that clearly signal shifts between parts (van Dijk and Kintsch; Kieras, "Initial Mention"; Frase). But apart from any natural disposition we may have to expect structure in text, our conception of text as an orderly succession of ideas is strongly reinforced by the constraints of the standard print medium: texts come to us on printed pages that we generally read in order, from the top down and from left to right.

Today, the constraints of the medium are being lifted by developments in computer technology. Visionaries of information technology foresee a time when most forms of written communication (from books and journals to reference manuals and mail) will be composed and

disseminated electronically rather than on paper. And instead of taking the traditional form of linear blocks of prose, such online material will be presented in hypertexts that link together individual bits of text and even whole documents.¹ Thus far, the most common application of hypertext has been for computer manuals, encyclopedias, or guide books, providing readers with immediate access to definitions of key terms, cross-references, graphic illustrations, or commentary from previous readers (Marchionini and Shneiderman; Yankelovich, Meyrowitz, and van Dam). If scholarly journals become routinely published in hypertext, readers may be able to move instantly from a citation in one article to the cited work or to any of the author's earlier or subsequent publications.

The advent of hypertext is a new and exciting development that has important implications for researchers and teachers in English. As Jef Raskin and others have noted, there is as yet a good deal of "hype" in hypertext and its full impact will not be felt in most English departments for a number of years. But the fact remains that sophisticated hypertext systems are increasingly available—commercially as well as within specific academic and nonacademic communities. Hypertext has the potential to change fundamentally how we write, how we read, how we teach these skills, and even how we conceive of text itself.

Hypertext promises to facilitate the writing process in several ways (Pea and Kurland). A writer's invention processes (generating and selecting ideas) may profit from opportunities to freely explore source material presented in a hypertext and make novel associations. The related processes of idea manipulation and organization, such as experimenting with various idea clusters or outlines, may be aided with a system that allows writers to create electronically linked "notecards" that can be sorted and rearranged (Neuwirth et al., "Notes Program"; Neuwirth et al., *Comments*; Trigg and Irish; Smith, Weiss, and Ferguson). Collaborative writing may be fostered by systems that enable peers to annotate each other's drafts or that help writers integrate individually written sections into a coherent draft (Irish and Trigg; Catlin, Bush, and Yankelovich). Hypertext systems may also be designed to meet specific pedagogical goals—for example, guiding novice writers through heuristic activities that support the critical thinking and analysis necessary to writing a policy argument (Neuwirth and Kaufer, "Role").

Apart from serving as a tool for writing, hypertext promises benefits to writers of computer manuals or reference materials. These writers typically face the problem of presenting large amounts of complex information to readers with wide-ranging needs—such as experienced

and novice computer users who may seek the same information but have quite different needs with respect to appropriate terminology, format, definitions, examples, and details. The task of these writers is further complicated by its subject matter, usually computer technology that changes even as writers scurry to describe it, so that printed material is outdated even as it is published—the “original sin” of computer documentation. The hypertext solution to these problems would replace printed manuals with an online network of information reflecting various levels of technicality (Robertson, McCracken, and Newell; Walker, “Authoring Tools” and “Document Examiner”). Readers with less technical expertise may choose to follow links to nodes with definitions, examples, explanations, reminders, or advice, which more sophisticated users may bypass completely. Or, instead of leaving the choices to the users themselves, hypertexts may be designed to guide readers on defined paths through the network at the appropriate level for their purpose or level of expertise (Zellweger; Younggren; Carlson, “Way”). A hypertext reference manual would ideally be suitable for all users, from novices to experts (and for novices whose skills develop over time), and for a variety of tasks. Such a system would presumably be easier to update than printed manuals and reduce the high costs of printing and reprinting.

Hypertext thus has a strong pragmatic appeal: to facilitate the efficient creation and dissemination of complex documents and sets of documents of all kinds and to allow people “to access information in the sequence, volume, and format that best suits their needs at the time” (Grice 22). The ultimate goal of these designers is to create a system so tailored to individual preferences and task situations that every user will feel as though entering an “information universe designed specifically for his [or her] needs” (Younggren 85).

In contrast to those who aim to micromanage the presentation of information in a network are the hypertext designers who are attracted to its Romantic side (Herrstrom and Massey). As Edward Barrett notes, “Developers of hypertext systems are inspired by a highly Romantic, Coleridgean concept of writing: an infinitely evolving text that tracks momentary cognitive processes within the individual reader-author” (*Text, Context, and Hypertext* xv). In this information age, hypertext Romantics aspire to a kind of unspoiled landscape of knowledge, dotted with visual and verbal outcroppings captured electronically. They view hypertext as a means to liberate readers (as well as writers) from the constraints of text boundaries, freeing them to wander through an array of connected texts, graphics, and commentary, to explore and create topical paths of associations at will. Such open-ended hypertexts

are being used in literature and other humanities courses to give students access to rich networks of cultural and historical material relevant to the primary texts under discussion (Beeman et al.; Dryden, in his response to this section).

Whether pragmatic or Romantic, the potential benefits of hypertext sketched here follow from certain assumptions about how people read or should read. The belief that readers can select for themselves which links in a network to follow rests on the assumption that readers know best what information they need and in what order they should read it. The goal of creating paths for different readers assumes that hypertext designer-writers can predict readers’ needs well enough to create the right set of paths and direct each reader onto the appropriate one. The very notion that hypertext designer-writers can create meaningful, useful networks in the first place depends on a whole range of assumptions about how to divide up and relate parts of texts, including which segments of text constitute meaningful nodes, which types of links are meaningful and important, and which types of texts can or ought to be read nonlinearly. In fact, many of these assumptions contradict current thinking in rhetorical theory, cognitive psychology, and document design. The evidence from these fields suggests that, as currently conceived, hypertext may dramatically increase the burdens on both readers and writers. My purpose in this essay is to review relevant educational and psychological research on reading that bears on the problems hypertexts may pose for readers and writers. My goal is not to accept or dismiss hypertext in principle but rather to point to specific aspects of reading and writing processes that hypertext designers must consider if they are to serve readers and writers effectively.

Obviously, readers approach hypertexts for a variety of reasons, from purposively seeking specific facts to browsing out of sheer curiosity (Slatin, “Reading”). I focus here on readers with more complex motives, those who read to learn, to understand and evaluate the ideas and arguments of others, to come to realizations about the subject matter, and to integrate what they have learned with what they already know. These goals will push hypertext to its logical extreme—the rich connection of the parts of a text to each other and to other texts. The hypertexts I have in mind are not so much online reference works or annotations of a single primary text (though the research reviewed here bears on those as well) as a fully interconnected electronic literature. While such hypertexts are, currently, the least well developed, they are also those most likely to influence what we in English studies do as teachers and as scholars.

Thinking, Learning, and the Organization of Memory

Many designers claim that hypertexts will facilitate reading and writing (and even thinking and learning in general) because, unlike linear texts, hypertexts closely resemble the networked, associational organization of information in human memory. This view probably originated with Vannevar Bush, who first conceived of hypertext, and has been carried forward in various forms to the present (Shneiderman, "Reflections"; Carlson, "Way"; Smith, Weiss, and Ferguson; Beeman et al.). While Bush's concept of human memory seems to have been quite advanced for his time, current hypertext proponents tend to misrepresent modern-day cognitive psychological perspectives on information processing.

The idea that hypertext is somehow more "natural" or more "intuitive" than linear text assumes a structural correspondence between networked information in a person's long-term memory and the presentation of information in hypertext network. This assumption contradicts some important, long-standing psychological findings about the organization of information in memory and the process by which new information is acquired.² First, although some researchers believe that information in memory is organized in completely amorphous associative networks, a great deal of knowledge seems to be arranged hierarchically and sequentially. Second, there is no evidence that readers can grasp information more easily or more fully when it is presented in a network rather than in hierarchical and linear form. The opposite may, in fact, be true. What people hear and see is not imported wholesale into long-term memory; it must first pass through a constraining "gateway." In particular, the processes of thinking and learning that draw on networks of previous knowledge are crucially constrained by the limitations of working memory (also referred to as *short-term memory* or *focal attention*).

Cognitive theorists posit working memory to account for human beings' inability to attend to more than a small number of things at any one time, regardless of whether these things are ideas recalled from prior knowledge or new information that has just been heard or seen or imagined. Further, what people attend to shifts over time; as they recall other ideas or observe new things, items that had been in focal attention "fade" or become "displaced" or "inactive." The shifting of attention imposes a kind of linearity or seriality on thought processes: since we cannot think about everything at once, we have to focus on a few things at a time in some order. A useful analogy for these ideas in long-term memory might be to imagine an auditorium

full of students. A variety of plausible principles might lead them to sit in certain groupings or might cause initially accidental groupings to take on significance over time. But the students did not enter in those configurations—their access to the room from the congestion in the hallway is constrained by a narrow doorway that forces them to enter in some sequence. The most efficient way to create an *intentional* configuration—one that facilitates taking attendance, for example, or that optimizes visibility for the group as a whole—may be to organize the students in advance, while they are out in the hall (as the marshals do at commencement ceremonies). Similarly, the fact that part of human memory may be arranged in associative networks does not mean that the best formats in which to read or write are also associative networks (Neuwirth and Kaufer, "Role"). If the goal is to ensure that readers consider a specific set of associations, then a highly organized text format is more likely to achieve that aim than an amorphous network.

The implications for hypertext can be stated even more directly. Because readers cannot import textual (or hypertextual) structures directly into long-term memory, the putative resemblance of hypertexts to long-term memory is irrelevant. The fact that hypertexts and long-term memory may both have networked structures in no way entails that hypertexts are superior to linear texts for facilitating reading or promoting learning. In fact, the development of linear text forms, with their careful sequencing of ideas, may not reflect constraints of the print medium so much as the needs of readers and writers who depend on the text to help them effectively sequence the flow of ideas through focal attention.³

Cognitive Models of Reading

A major premise of most reading theories, consistently supported by empirical studies, is that, as people read, they build a hierarchically structured mental representation of the information in the text (Kintsch and van Dijk; van Dijk and Kintsch; Meyer; Just and Carpenter). As they read successive sentences, they link the ideas or propositions expressed in them to their developing hierarchical representation by means of chains of repeated concepts (or arguments). To the extent that the sentences—or larger units—of the text reuse, develop, elaborate on, and interrelate the same arguments, the text is more cohesive. The more cohesive the text, the easier it is for readers to create a well-structured, meaningful, and useful mental representation (Eylon and Reif). The quality of the representation, and the ease with which it is constructed, crucially depend on the order in which readers encounter

the propositions and on the amount of repetition and development of important concepts (or "arguments") in successive portions of the text. It is more difficult to create a mental representation of a disjointed or disorganized text. If readers come to a sentence that seems to contain no previously encountered arguments—that is, has no obvious link to the surrounding sentences—they must either retrieve from memory earlier propositions that contain one or more of the arguments or infer some link between the sentence and some part of their representation of the text. Both retrieval and inferencing are relatively costly processes in terms of time and effort. Working from the assumption that only propositions currently "active" in working memory can be linked, researchers have successfully predicted which kinds of texts are easier to read, understand, and remember than others. Bruce Britton and Sami Gülgöz have recently gone even further, using Walter Kintsch and Teun van Dijk's model to identify sites for textual revisions that resulted in improved comprehension and recall.

As Catherine Smith notes, in chapter 11 of this volume, Kintsch's model has evolved over the years. In its original form (Kintsch and van Dijk), it was fairly rigid and deterministic—allowing little into the mental representation of a text beyond literal decompositions of the sentences and necessary bridging inferences. It did, however, successfully predict what parts of a text are best remembered. In later work (van Dijk and Kintsch), Kintsch fleshed out various parts of the model to account for other kinds of information that readers call on regularly (such as knowledge of genre and situational knowledge). His current position ("Role") extends the model still further, allowing for many more idiosyncratic (and even inappropriate) associations to end up in the mental representation of a text (but also requiring an additional "cleanup" process). This position in no way represents a recantation of his earlier views, as David Dobrin suggests in his response to this section. To the contrary, Kintsch explicitly retains most aspects of his model—and, in particular, the ones relevant here—frequently referring readers to his previous work for elaboration of them (see esp. 166, 167, 168, and 180). In essence, Kintsch's innovations enrich the mental representation or "text base" by allowing more associations to the reader's general knowledge. They do not change the effect of the text itself on the construction of the text base and therefore do not lessen the importance of beginning with a coherent, well-structured text. In Kintsch's "revised" model, as in his original model, the sequence of sentences and sections of a text and the explicitness of their connection to one another largely determine how well and how easily a reader can construct a text base.

Several text features have been identified that consistently make it easier for readers to construct a coherent representation of a text, to reflect on its relation to prior knowledge, and to integrate new ideas and new information with what they already know (Felker et al.; Kieras and Dechert). First, for readers to make appropriate connections between related ideas, the sentences expressing these ideas should appear in close proximity. Thus a text is easier to read if its points are developed in coherent sequences of sentences, paragraphs, and sections and if it contains discourse cues that signal the relations among these ideas (Halliday and Hasan; Fahnestock; Britton and Gülgöz). Second, since readers use high-level ideas to tie portions of the text together, these concepts should be explicitly stated early in the text and should be clearly signaled so that readers can easily recall them as the need arises (Kieras, "Initial Mention" and "Model"). Thus it is easier to read, comprehend, and remember a text if it contains an informative title, headings, overviews, and topic sentences introducing key concepts that are repeated and developed in successive portions of text (Schwarz and Flammer; Glynn and Di Vesta; Mayer, Cook, and Dyck; Wilhite). Reading is also easier when the text reminds readers of relevant points (normally through repetition or reference to the earlier discussion). Finally, while readers are capable of following innovative text structures (especially when the text announces its structure explicitly), the easiest texts to read are those based on a familiar structural pattern or genre (Meyer and Freedle; van Dijk and Kintsch).

The strategies for structuring texts described here are not unfamiliar ones. They are the product of centuries of experimentation by writers striving to make their texts more comprehensible to readers. These strategies, however, place the burden of selecting and arranging information, and providing signals to the arrangement, primarily on the writer. Hypertexts, by shifting a large portion of this burden to the reader, by proliferating the readers' choices about what portions of a text to read and in what order, compound the difficulties of creating a coherent mental representation.⁴

Effects of Text Structure on Reading and Learning

Many reading theorists believe that after reading a number of texts with similar structures, such as a series of fairy tales, newspaper articles, or research reports, people formulate generalized, abstract patterns or frameworks, called "schemas," that they call on as they encounter new texts of the same type. As they realize they are reading a familiar type of text, they invoke their schema for that genre and use it to anticipate

what will occur next, to make inferences to fill in implicit or missing elements, and, later, to reconstruct the text from partial memories. People often rely on the structure of the text and the expectations raised by schemas to decide which aspects of the text are most important and, accordingly, where to allocate their time and attention during reading (Just and Carpenter). Further, once a schema is invoked, information in the text that fits the pattern is integrated easily, but information that seems peripheral or incongruous tends to drop out—either it is never linked to the mental representation of the text or, if it is encoded, the link to it is so weak that it effectively is lost (Bartlett).

Schemas have been posited for other cognitive processes besides reading. Many cognitive theories assume that much of the knowledge in long-term memory is organized around such hierarchical frameworks (referred to in various theories as *schemas*, *frames*, or *scripts*) that capture familiar patterns among elements. There may be schemas for events, for genres of text, for characteristics of a species, for the elements in a system. As Smith points out in chapter 11 of this volume, Kintsch (“Role”)—along with other psychologists—has come to reject the schema as a cognitive *mechanism*, that is, as a way to formalize or model the way in which encountering a familiar proposition reliably evokes a pattern of related propositions. Neither Kintsch nor other psychologists, however, will dispute the consistently observed *behaviors* that schemas are meant to capture.⁵ Regardless of what cognitive mechanism is ultimately selected as the best formalism for the phenomenon, the concept of a script or schema remains a useful one.

Readers invoke a particular schema in part because of cues provided early in the text, such as the title or the initial sentences. The remainder of the text then may either fulfill the expectations raised by the schema or confound them. When texts set incoherent expectations or fail to confirm expectations they initially raise, they create problems for readers, especially those to whom the subject matter (or “domain of knowledge”) is unfamiliar. Bonnie Meyer studied this problem by creating texts that raised expectations for one structure (e.g., problem-solution) but actually developed according to another (e.g., comparison-contrast). She studied how well readers coped with such texts, including both readers who were familiar and readers who were unfamiliar with the subject matter (“domain experts” and “novices”). She found that the novices relied heavily on the text’s structure to create their representations and were therefore misled by the opening portions of the text. In contrast, the domain experts were generally able to recover from the textual miscues and construct coherent representations. Experts can draw on their knowledge of the domain’s concepts and principles

to determine the centrality or novelty of textual information, regardless of where it appears in the text structure.

The structure, or organization, of a text thus signals the relative importance of its various parts, influencing how readers allocate their time and attention and thereby influencing what information they are likeliest to remember. But the way a text is organized can also influence the textual effects that Joseph Grimes has termed “staging”—how easy it is for readers to reflect on the ideas they have read, to juggle and compare them, to see how those ideas relate to one another and to other ideas they have on the subject. We can illustrate the demands that different organizations impose on readers by considering the familiar problem of how to organize a comparison-contrast essay. The two most common strategies, which appear in scores of writing textbooks, are to organize (1) around the objects or alternatives being compared or (2) around the points of comparison—that is, around various aspects of the objects being contrasted or the criteria against which alternatives are judged. In his excellent technical writing textbook, Paul Anderson schematically represents these patterns (266). The first strategy, the “divided pattern,” uses the objects as superordinate terms, repeating the criteria (in a consistent sequence) under each object heading. The second strategy, the “alternating pattern,” sets up the criteria or aspects as superordinate terms, under which the discussion alternates among the objects under analysis (again in a consistent sequence for each aspect).

The choice between these strategies, of course, has rhetorical implications. For example, an aspectual orientation may be more appropriate for a technical feasibility report whose varied audiences may each have a specific interest in one criterion or another (e.g., cost, efficiency, environmental impact). But the choice also has important implications for readability as well; the two organizations impose different burdens on comprehension processes. Wolfgang Schnotz (“How Do Different Readers Learn” and “Comparative Instructional Text Organization”) argues that the aspect-oriented (alternating) organization is the more difficult because readers must switch attention back and forth between different objects; each switch requires reactivation of the reader’s prior knowledge and current representation of the object. This switching is especially difficult when the reader is unfamiliar with the topic.

Building on the Kintsch and van Dijk model of reading comprehension, Schnotz maintains that because the two organizations put different propositions into close proximity, readers create different representations of their content in memory. In particular, readers of

an object-oriented (divided) organization are more likely to create a well-integrated representation of each object but will find it more difficult to keep track of their similarities and differences. The text itself does little to push the reader to form these interconnections (though the reader is, of course, free to do so). In contrast, readers of an aspect-oriented organization focus on the similarities and differences and as a result are also forced to develop cross-referenced representations of each object (through the costly switching process). Schnotz studied how these organizations influence a reader's ability to recall the overall meaning of the text and make accurate comparisons between the objects. He found, as expected, that aspect-oriented text took longer to read, but readers of these texts remembered more and were better able to make sophisticated discriminations between the objects. Further, the readers' familiarity with the topics significantly affected their ability to cope with the two organizations. Readers with little previous knowledge learned more from the object-oriented texts than aspect-oriented ones, presumably because they could avoid switching. Conversely, readers with more previous knowledge learned more from aspect-oriented text; they were able to take advantage of the close proximity of the comparisons across objects.

These studies highlight the importance of the order in which readers see information. Hypertexts, which proliferate the possible sequences, raise significant issues for both readers and writers. For example, it is easy to imagine a hypertext version of a comparison-contrast essay that allows readers to choose an aspect-oriented or object-oriented organization. What choices will readers make? Will those with little domain knowledge realize that an object-oriented organization will be easier to read? Will they be aware that working through the aspect-oriented organization will be worth the effort for learning careful discriminations?

Implications of Cognitive Models of Reading for Hypertext

Because of the cognitive view of the reading process just described, it is easy to see what potential problems hypertext may raise for readers, for the very reason that hypertext violates standard assumptions of what texts are like. Readers traditionally rely on the writer to select topics, determine their sequence, and signal relations between them by employing conventional discourse cues. The net effect of hypertext systems is to give readers much greater control over the information they read and the sequence in which they read it. Along with greater control, of course, comes a greater burden for the readers, who must

now locate the information they need and relate it to other facts in the network, often without the aid of traditional structures or discourse cues.⁶

Many hypertext designers recognize the problems such networks may present, especially for readers who are unfamiliar with the concepts in the text. They report informal evidence that users become overwhelmed by the choices among links and by the difficulties of maneuvering through the networked text structure (Conklin, "Hypertext"). As a result, readers lose track of where they are in the network (and where they have been) and often read a great deal of material that is irrelevant to their purpose (Foss; Yankelovich, Meyrowitz, and van Dam; Whiteside et al.). Technologies related to hypertext have also been shown to pose significant problems for users. Stephen Kerr cites one study on a menu-selection Videotex system in which 28% of users gave up without finding the information they wanted even though they knew that "the information was in there somewhere" (333). Half the users of another menu-driven system had to backtrack at least once before finding the information they sought. While recognizing the navigational difficulties in general terms, hypertext designers have not weighed some of their deeper implications for reading and writing processes.

Consider the hypothetical case of readers using an open-ended hypertext, one without predefined paths, who must choose what links to follow through a set of connected texts (each of which is also represented as a network of nodes). Assume that the goal is not to read everything in the network but instead to gather information relevant to some particular issue. First, since readers must choose what to look at, they may never see all the "right" information, either because they cannot find it or because, for some reason, they fail to select it. Second, even if they do see the "right" information, they may see it at the wrong time. As described above, the timing of seeing a particular bit of information could determine whether readers judge it to be important or whether they see its connection to information they have already read or have yet to read. If they do not see it in conjunction with other relevant information, they may have to expend great effort to integrate it coherently into their mental representations. If they fail to do so, they are likely quite literally to miss the point. Third, readers may see a great deal of intrusive, irrelevant information that may skew their representations. Even if they recognize that some information they have read is irrelevant, there may still be adverse consequences of having spent time reading it. Finally, readers may lose a sense of the integrity of any given text in the network, since they may be unaware

of crossing from one text to another. Lacking a sense of textual integrity, they may have difficulty relocating information they have read or attributing the ideas to the correct sources. In short, in addition to suffering the frustrations of disorientation or cognitive overload that hypertext designers already acknowledge, readers may come away with a false or incomplete representation of the texts in the network or even the information relevant to their topics.

The worst-case scenario sketched here is speculative; little research has been conducted of the actual effect of hypertext on reading. However, available research, some using printed texts and some online materials, addresses the specific issues raised in this case. Can readers make appropriate selections of what and how much to read? Can readers create appropriate sequences of textual material? If readers are unable to navigate a hypertext effectively, can hypertext designer-writers reasonably anticipate readers' various needs and create appropriate paths to satisfy them?⁷

Can Readers Select What and How Much to Read?

Many hypertext designers assume that readers know what sequence of information is best for them, that they can tell when they have read enough or judge whether what they are reading is important. However, the evidence suggests that readers are not very good at assessing the adequacy of the information they have encountered and are even worse at anticipating whether important or useful material remains in the portions of text they have not reached.

David Kieras ("Role") found that many readers, left to decide how much to read, stop too soon. In his study, adults with varying technical backgrounds were given online, step-by-step instructions for using a mechanical device. The instructions were presented in a hierarchical network that organized the steps according to major tasks; the bottom-most level contained the directives for specific operations, such as turning on a specific switch. At any level of the hierarchy, participants had the option of reading on to a deeper level of detail (using a menu-selection system) or attempting to carry out the steps. Kieras found that the participants tended to stop reading before discovering crucial details—presumably with the impression that they understood what to do—and, as a result, failed to carry out the instructions correctly. In contrast, participants who read and followed the instructions presented in traditional linear order were much more successful at completing the task.

David Reinking and Robert Schreiner also found evidence that readers may fail to take full advantage of useful information available to them in hypertext. In their study, fifth- and sixth-grade students were presented with a set of expository passages annotated with various online aids such as definitions, paraphrases, background information, and distilled main ideas. Students who were allowed to select at will from these aids performed significantly worse on various comprehension tests than did students who were guided through all the aids. In fact, the free-selection group performed more poorly than did students who read the printed version of the linear text without any aids.

While Reinking and Schreiner's results are based on the activities of schoolchildren, they are consistent with Kieras's findings with adults. Taken together, these studies suggest that when readers are responsible for selecting what text to read, they often omit significant information altogether, perhaps because they can't find it, they don't know it's there, or they don't think it's important.

Can Readers Create an Appropriate Reading Order?

The view that readers can select for themselves which links in a hypertext to follow is based on the belief that readers know best what information they need and in what order they should read it. Hypertext designers thus assume that readers can organize information appropriately for their level of knowledge and their purpose in reading. In fact, little research has been conducted on how readers themselves choose to sequence the pieces of a text, whether reader-chosen orders are generally different from those a writer or teacher might create, and what effects these different orders have on what readers learn. As the following discussion indicates, the available evidence is mixed. Overall, it suggests that certain kinds of readers, or readers in certain situations, may benefit from the active effort required to sequence reading material for themselves.

Sequencing an Entire Network

Hans Lodewijks evaluated a variety of text sequencing systems, or "presentation orders," using printed materials. The materials consisted of sixteen passages written for high school students on concepts from electricity (e.g., ampere, conductivity, electron) and presented to different groups of students in different sequences. The "teacher-regulated" sequences were similar to guided paths through a network, giving

students little control over the order of the passages. They included a sequence determined intuitively by a group of physics teachers, others based on various logical dependencies or cross-references among the concepts, and an alphabetic sequence based on the concept headings. Two "self-regulated" systems allowed students to choose a concept to study from an alphabetical list, read the passage, and then return to the list to select the next concept. One self-regulated system also provided a structural overview of the relations between the concepts, similar in some ways to the graphical maps (or "browsers") provided in some hypertext systems. In general, the self-regulated sequences (especially self-regulation with a structural overview) led to better recall and better recognition of relations and inferences among the concepts than any of the teacher-provided sequences.

However, not everyone benefited from self-regulation. In particular, Lodewijks found that "field-dependent" learners and low scorers on various logical reasoning tests performed poorly under self-regulation conditions but significantly better with teacher-regulated sequences. The converse was true for "field-independent" learners and high scorers on reasoning tests: they performed significantly better with self-regulated texts than with teacher-regulated ones.⁸ Thus the readers' preferred learning strategies (or "cognitive styles") may determine how well they can cope with charting their own path through a hypertext.

Richard Mayer conducted a similar study using instructional materials for writing computer programs, presented as a set of printed cards. He investigated how "experimenter-controlled" and "subject-controlled" card sequences influenced the ability of college students to solve programming problems. In the experimenter-controlled sequence group, students read the cards in either a logical order or a random order. Students in the subject-controlled group were given a table of contents for the cards, in which the topics were listed either in random or in logical order. They used the table of contents to pick which card to read next. Mayer found no overall differences between the sequences. However, he did find differences in the types of problems participants were able to solve. Participants who had chosen their own reading order were significantly better at solving novel, unexpected types of programming problems, while those who had read the text in experimenter-controlled sequences were significantly better at solving problems similar to those in the text. Mayer believes that allowing people to choose their own reading order "may result in deeper, more active encoding, which allowed subjects to struggle harder to relate the text to their own experience rather than memorize the information as presented" (149).

Both Lodewijks and Mayer used printed materials that simulated the conditions of hypertext in many important respects. Two researchers have recently conducted studies of reading in hypertext itself; the results are largely consistent with those based on printed materials.

Sallie Gordon and her colleagues ("Effects of Hypertext") found that reading in hypertext format decreased information retention as compared to studying an online linear presentation. Gordon et al. constructed hypertexts for four expository texts—two on technical topics and two on general-interest topics—each about one thousand words long. In the hypertexts, main ideas were kept on the topmost level and links to elaborative text segments (such as examples, definitions, and noncentral information) were assigned to a second or sometimes a third structural level (about half the text was presented in deeper levels). A highlighted keyword in a main-level segment signaled the presence of deeper information, which readers were free to access by pressing a key. The participants, upper-level college students, were told to use normal, casual reading processes for the general-interest articles but to study the technical texts carefully. After reading both a linear and a nonlinear text, the students were asked to recall as much of both texts as possible, to answer questions about them, and to express their preference for reading format. Gordon et al. found that for both types of texts, students who read in the linear format remembered more of the basic ideas and, for the general-interest articles, assimilated more of the text's macrostructure than after reading in hypertext. Most students also preferred the linear presentation, perceiving it as requiring less mental effort. As a result of finding such negative results for hypertext, Gordon and Vicki Lewis have sought more effective ways to segment material in order to create more easily processed hypertext structures ("Knowledge Engineering" and "Enhancing Hypertext").

Jean-François Rouet ("Interactive Text Processing") found that sixth- and eighth-grade students have difficulty making appropriate sequencing selections with hypertext materials. Rouet constructed hypertexts for four general-knowledge domains. A hypertext consisted of six related thematic units, each containing a title and a fifty- to seventy-word paragraph—representing a "chained list rather than a network" (253). The hypertexts differed in the availability of various cuing aids, such as markers for previously read topics, availability of the topic menu during reading, and explicitness of statements relating one topic to another. Students were asked to read all of a hypertext, selecting the topics in any order and as often as they wished. Then students answered multiple-choice comprehension questions and wrote a summary. Each student read all four hypertexts, two in the first session

and two a week later. Rouet computed various measures of selection efficiency, including the number of repeated readings of a topic (indicating global orientation difficulties) and the number of times students picked illogical sequences of topics (indicating local orientation difficulties). Although grade level and some combinations of cuing aids improved performance, Rouet found evidence of global and local disorientation at both grade levels, even with his very simple nonlinear structures. For example, in only 35% of their selections did students pick the topic that was most closely related to the one they had just read. Explicitly marking relations between topics improved students' appropriate selections only to about 50%. Practice at using the system evidently also helped somewhat; the percentage of appropriate selections in the second session increased to 58%. We should not conclude from this study, however, that students can learn to cope with any hypertext with practice. These students may have eventually figured out the simple and consistent structure of these hypertexts, especially when aided by explicit textual cues. Accordingly, these results may just mean that students can improve somewhat at using well-marked and structurally predictable networks.

Obviously, these studies have several important implications. Rouet's study indicates that students may have difficulty making their way through even simple hypertexts. The work of Gordon and her colleagues suggests that reading from hypertext can actually impede a student's comprehension of a text, relative to a linear presentation. However, Lodewijks's and Mayer's investigations hold out the promise that at least some students (those with particular cognitive styles or reasoning ability) may learn more effectively when they choose their own reading order instead of following sequences imposed on them by teachers or writers. Further, self-regulation forces readers to adopt more active reading strategies, which generally lead to better learning.

Several circumstances limit our ability to draw clear-cut conclusions from these studies. First, unlike most hypertext reading situations, participants knew that they had to read the whole text, that everything they needed to learn for the test was in the network, and that they had to learn all of it, to make it all fit together. Second, in all four studies, participants used a finite list to select each successive topic. In the context of reading to understand the network as a whole, the task of ordering the segments reduces to a puzzle: looking at the topics left on the list and guessing which one would be best to read next. This is quite a different task from selectively browsing through a large, messy network. Third, only in the case of Gordon et al. was the text based on an integrated piece of prose; in the other studies, the material was

developed from individual modules on fairly discrete concepts. Fourth, neither Lodewijks nor Mayer described the orders students actually devised. Thus, it is unclear whether the benefits of self-selected orders resulted from some feature of the orders per se or from the very fact that the readers were forced to think about how to sequence the text. That is, were the teacher-regulated orders deficient in some way or do learners simply benefit from actively puzzling out how to arrange the text and make sense of it?

Overall, then, these investigations created conditions that encouraged the "self-regulating" students to seek actively for the connections among a finite set of textual elements. Some students in Lodewijks's experiment thrived under these conditions, but not all of them. A task like this may actually encourage students to read more actively (though it's not yet certain that the benefits of these exercises persist). However, for most purposes, readers are unlikely to devote the time and energy necessary to fit all the pieces of a network together, and hypertext designers may not be as fastidious as these researchers about selecting appropriate information to include in the hypertext, to provide explicit relational cues, or to create relatively simple and predictable structures.

Sequencing Selected Information from a Network

Another important situation to consider occurs when readers must select portions instead of having to read the whole text. Carolyn Foss investigated the effects of self-regulation within a large hypertext in which readers controlled both what to read and when to read it. Her participants were asked to compare and analyze information distributed across a hypertext network (Xerox Notecards) to solve a specific problem. The network consisted of encyclopedia entries for a set of ten countries that were not identified by name. Facts about various aspects of each country (population, climate, etc.) were available on "cards" that the user could pull out of a "file," arrange on the screen, or refile. Given a list of the countries' names, the users had to read and compare facts about the countries in order to guess the identity of as many countries as possible within a set time period. These users, then, had to select which cards to read, which ones to leave open, and how to arrange them.

Foss's participants varied greatly in searching skills and the ability to manage the clutter of open entries on the screen. About one-third of them opened very few entries at any one time—and, accordingly, kept their displays rather neat. These users read more total entries than

other participants but were unable to make effective comparisons. They identified few countries correctly and took longer to do so. The remaining participants, who kept many entries open at once, were more successful at the task, presumably because they could view more information at the same time. About half of them followed systematic search strategies and used the screen display efficiently. The rest were highly unsystematic, were easily sidetracked, and wasted time revisiting cards and sorting through the cluttered display.⁹

Foss's study reinforces the notion raised earlier that it is crucial for readers to see relevant information in close proximity in order to make appropriate connections. The participants in this study who kept a large number of cards open were more successful at this particular task because they were able to see enough facts at one time to make useful comparisons and notice useful details. The participants who were more worried about keeping neat displays may have read exactly the same cards but failed to make the connections—presumably because they read them at an inopportune time or in fruitless conjunctions.

It follows, as a more important implication of Foss's study, that many people are not very good at regulating their selection and organization of information. Of the participants in the study, fully one-third adopted unproductive "neat-screen" strategies. Another third created wildly messy screens; this strategy seems not to have cost them much in this task but may well create problems in other tasks (just as a neat-screen strategy might be more advantageous in another kind of task). It is unclear, of course, whether people maintain their preferences for neat or messy screens in different task situations, or whether people who use hypertext more often learn strategies appropriate to the task at hand.

In addition to individual preferences for how to manage a display, differences in previous knowledge of the information in a hypertext influences selection strategies. Rouet ("Initial Domain Knowledge") investigated how prior knowledge affected the efficiency of high school students looking for information in a hypertext network in order to answer specific questions. He found that students who were highly knowledgeable about a particular subject area were much more efficient at locating the relevant information than were those with low or moderate knowledge. Furthermore, by showing that students' performance declined when they moved from a familiar to an unfamiliar domain, Rouet demonstrated that the effects resulted from specific domain knowledge and not from general reading ability or practice with the hypertext system. Taken together with the studies by Meyer and Schnotz described earlier in this chapter, Rouet's findings under-

score the hardships that face a hypertext user who attempts to learn about an unfamiliar domain.

These studies suggest that the best way to sequence information is not at all obvious to readers. If the goal of hypertext is to allow people "to access information in the sequence, volume, and format that best suits their needs at the time" (Grice 22), the results of these studies indicate that a large proportion of hypertext users will need a good deal of guidance in determining the most appropriate sequence, volume, and format of information.

Can Hypertext Designers Create Appropriate Paths for Readers?

Anticipating that readers will have trouble charting a logical course through a network on their own, some hypertext designers are designing alternative paths through a network to which readers would be directed as appropriate (Younggren; Zellweger). For example, Patricia Carlson ("Way") describes how the reader's task may be used to organize the technical documentation for an airplane into several orders. An alphabetical order would be most appropriate for someone trying to update selected pieces of information. An order based on spatial layout would be suitable for someone who had to work efficiently from one end of the plane to the other, performing a set of repairs. Alternatively, the most efficient order might be to group information according to which tools are needed for maintenance or repair. This solution assumes that designers can anticipate the necessary paths and can reliably determine which readers need which path.

Creating such paths is hardly a straightforward task. Various factors, including the designer's own idiosyncrasies and sensitivity to readers' needs, the type of text being incorporated into the network, and the choices that the hypertext system itself allows, influence what kinds of sequences any given hypertext designer may create. The best evidence of the difficulties involved comes from an attempt, by three leading hypertext designers, to use their own systems to construct hypertexts of one small set of documents, six articles from the proceedings of a hypertext conference. Liora Alschuler carefully analyzed the resulting hypertexts and found vast qualitative and quantitative incongruities in how the articles were segmented and linked, both within any one system and across systems. Alschuler attributed one source of difficulty to the designers' inconsistency and their, perhaps inevitable, subjectivity. However, a second stumbling block was the nature of the hypertext

systems themselves, which imposed significant constraints on the kinds of segments and links that were even possible. This uncontrolled variability leads to serious methodological problems for those researching the effect of hypertext on reading and writing: How can one draw generalizations about hypertexts if no two systems produce the same result from a set of source texts and even the texts in any one given system are subject to erratic treatment?

I am skeptical that a hypertext designer, even under ideal conditions, can anticipate all the paths that readers may wish to create within and between texts. As we have seen, a wide range of factors influence the appropriateness of a sequence for a given reader, including the reader's prior knowledge of the domain, the reader's task or purpose for reading, the reader's learning style, and the nature of the information itself. Because of the huge number of possible combinations of such factors, the array of alternative paths that a designer might create becomes a practical impossibility—and there still remains the problem of directing the right readers to the right paths. Obviously, these are issues that only experimentation with hypertexts can resolve; my purpose here is to raise serious concerns that hypertext developers should take into account, not to discourage the development of hypertexts altogether.

Others have raised the issue of what kinds of texts can be or should be integrated into hypertext networks. Many agree that a text with closely interwoven points is not an easy or desirable candidate for conversion to hypertext because "it destroys the subtle interconnections of theme, argument, metaphor, and word choice" (Carlson, "Intelligent Interfaces" 63; Shneiderman, "Reflections"). Darrell Raymond and Frank Tompa note that, because converting a text into hypertext makes implicit structures explicit, "the key question in conversion must be *will explicit structure be as expressive as implicit structure?* When the answer is yes, the document will gain from conversion; otherwise, conversion will degrade the representation of the document" (146). Despite the sense that texts with complex internal structures (including many forms of narrative, expository, and persuasive prose) may suffer from conversion to hypertext, they are the very sorts of texts that many designers want most to include.

Implications of Cognitive Processes for Hypertext Design

The evidence from cognitive psychology reviewed in this essay emphasizes how heavily we rely, in dealing with the world, on systematic

patterns of information. Our dependence on predictable patterns creates an enormous tension between the impulse toward creativity, inventiveness, and imagination, and the more conservative, "normalizing" forces that assimilate new information to established and familiar patterns. The cognitive mechanism for encoding information in long-term memory (a process that requires sustained or repeated conscious attention) and the selectivity imposed by schemas (i.e., the loss of things that don't easily fit) are both strongly conservative forces. It seems reasonable that we have mental mechanisms such as these, given the constant barrage of observations and sensations presented to us at every moment from our senses, our emotions, our intellect. Such mechanisms may account for our ability to "make sense," to impose order on the world.

The Romantic view of hypertext that aims at enabling imaginative leaps and connections between disparate texts, facts, and images thus puts enormous technological and creative effort at the service of preserving what might be quite rare and ephemeral associations. Some of these connections, probably a very small proportion, may be of great value and interest to those who initially make them. However, once the insight or connection is made, it is unclear that the thinker needs or wants to store the convoluted trail of associations that led up to them, let alone those that led nowhere. Such trails are probably of even lesser value to subsequent readers. The trails of associations in a hypertext may represent the ultimate in what Linda Flower calls "writer-based prose," prose that reflects the writer's process of coming to terms with a set of ideas but that may bear little relation to his or her final stance and none whatsoever to the readers' needs. Some proponents (e.g., Beeman et al.) claim that allowing students to explore freely in hypertext may foster insights and critical thinking through the creative juxtaposition of ideas from multiple perspectives. However, instead of prompting truly original insights, this process may simply reduce itself to a guessing game, as the user figures out what the hypertext writer (usually the teacher or another student) had in mind when creating a link. As discussed earlier, the nature of the reading process suggests that chance conjunctions and odd juxtapositions tend to be dropped from the reader's mental representation of a text. Thus the most imaginative links that a hypertext writer-designer creates are unlikely to be remembered or to influence the reader's subsequent thinking in any significant way. But, for the reader, the consequences of whimsically following links into disparate texts or text segments may be to obstruct or hinder the more conventional but durable processes of systematically integrating new information with old.¹⁰

To the extent that readers rely on structure—that they learn by discerning and internalizing the structure of a text (Jaynes)—hypertexts that emphasize free-form browsing may interfere with readers' efforts to make sense of the text and even with more limited and pragmatic efforts to find information relevant to some specific question. A hypertext system that promotes free exploration and browsing, then, may be effective only in certain kinds of situations, such as reading for pleasure but not for scholarship. In contrast, a hypertext system that allows readers to choose only among a fixed set of paths through the network may satisfy particular readers' purposes better—though designers of these systems face significant challenges for creating the right paths and steering the right readers onto them. In these systems, though, the romance of hypertext disappears; the hypertext becomes functionally identical to a set of linear texts.

As part of a literate society, we are familiar with traditional text structures. We have time-tested cognitive and rhetorical theories to bring to bear on describing effective printed texts, and we have derived from these theories a wealth of practical advice to convey to writers—students and professionals alike. But we lack corresponding theories for how to deal with hypertexts—especially those that push the limits with complex linkages within and between a complex set of texts. In this essay, I have sketched the challenges that an effective hypertext would face. Much work remains to fulfill the promise of hypertext for readers and writers.

Some designers are attempting to overcome problems of disorientation and cognitive overload (e.g., Utting and Yankelovich; Rouet, "Interactive Text Processing"; Gordon et al.; Gordon and Lewis), but more research is needed. The critical issues fall into two broad categories: the construction of hypertexts and the effects of hypertext on the reading process. The first order of business may be to create, and to study the creation of, large, complex, sophisticated hypertexts, involving a range of texts. As in Foss's study, some of these hypertexts should contain expository information intended to help people solve research problems. Others should be constructed to explore the effect of hypertext on existing imaginative literature, to supplement the work of "interactive fiction" writers who are designing literary works specifically for the hypertext environment (e.g., Moulthrop, "Hypertext and 'the Hyper-real'"). Finally, in spite of the fact that hypertexts were originally intended as resources for scholars, few complex networks yet exist for scholarly literatures, such as books and journal articles. These studies on the construction of hypertexts should, like Alschuler's, investigate the constraints imposed by different hypertext technologies. Drawing

on rhetorical, linguistic, psychological, and literary theories of text structure, these investigations should confront the issue of what kinds of texts can or should be presented in hypertext environments and, for those that are appropriate, should systematically explore different ways to partition and interrelate a set of texts. Formative evaluation of such hypertexts with readers should be ongoing. Research encompassing the full range of factors identified in the research reviewed here must eventually involve readers with different purposes, cognitive styles, amounts of background knowledge. The goal of such research must be to find ways within hypertext to provide appropriate discourse cues, cues that help readers decide what to read, how much to read, and when to read the rich array of information available in the network. These signals may be manifested in hypertext design in a variety of ways—many of which have been inadequately explored in current systems (as discussed by Wright and Lickorish). Hypertexts are here to stay, but it is up to researchers, teachers, and software designers to ensure that these texts promote the work of writers and readers.¹¹

NOTES

¹There is an extensive and growing literature on hypertexts. For a history and overview, see Conklin, "Hypertext." For bibliographies, see Mitterer, Oland, and Schankula; Harpold, "Hypertext."

²For a straightforward, general introduction to cognitive psychological perspectives on memory and learning, see Bransford. Many of the issues raised in this chapter are also discussed by Just and Carpenter and by Sanford and Garrod, who all focus on those aspects of cognition that relate to reading. For a brief review of the history of cognitive psychology, see Hayes.

³Some hypertext proponents assume that these cognitive limitations somehow result from the limited forms of writing available through print technology. They believe that a revolution in technology will lead to a revolution in cognitive capacity—an analogy with the advances presumed to have taken place when oral cultures became literate. In fact, many researchers have concluded that literacy did not change human cognitive capacity but, rather, the way that capacity is used (for general discussions, see Borland; Rose, "Narrowing the Mind"). In particular, technology is unlikely to influence the way in which we visually process written texts or auditorially process speech, both of which are linear (for reviews of relevant research, see Just and Carpenter). We read in linear sequences of words and sentences. While hypertext may change which sequences are available—and may well impose more frequent decision points for which sequence to follow—it will not change our basic mental architecture.

⁴The shift in burden also has consequences for writers. Carlson ("Way") raises the issue of whether and how writers can create segments of text that will make sense to any reader who encounters them, without knowing what previous segments he or she has read: "Since nodes may be 'threaded' into different paths, they must be reusable, requiring research into the requisites (as well as the effects) of free-standing, rhetorically 'neutral' prose" (96). The concept of "rhetorically 'neutral' prose" is one that many rhetoricians find troubling.

⁵In fact, the mechanism that Kintsch adopts, spreading activation within a network of links of various strengths, in effect functionally reproduces the evocation of a schema. The mechanism works by activating a set of propositions that are most strongly linked to some other proposition—presumably including many previously assumed to reside in a schema—while allowing for the activation of additional, less predictable kinds of knowledge and associations. It is evident from the recent research literature that cognitive psychologists do not take Kintsch's current position as rendering the terms *script* or *schema* meaningless—researchers who cite and build on Kintsch's 1988 article ("Role") continue to explore how schemas and scripts influence learning and memory (e.g., Maki).

⁶A more familiar and analogous case of the increased burdens imposed by technological "liberation" is the advent of desktop publishing. Such systems do indeed free writers to make their own choices concerning details of typography, page layout, graphics, and the like. However, the price of greater control to individual writers is the responsibility (and sometimes the sheer necessity) of acquiring some expertise in effective graphic design and printing.

⁷Some of these problems, of course, apply to conventional printed text too. People don't have to (and don't) obediently read in order (Charney, "Study"). Active readers often break out of a linear sequence to seek an overview of the paper—perhaps by reading the abstract, introduction, and conclusion and then skimming over the headings and looking at diagrams and tables. Depending on their purpose, they may skim for the main points of each paragraph or read only sections relevant to their goal. The point is that hypertext makes the situation much worse. Readers of printed texts can scan the page or thumb through the text to see how much is there, to see whether anything further down the page looks important or useful, or to see how much elaboration the writer considers it worthwhile to include. Obtaining an overview is a lot harder to do in hypertext because readers must consciously choose to look at something (i.e., by clicking a button), because new information may displace other information on the screen, and because making such a choice may make it harder to get back to where the reader started.

⁸*Field dependence* and *field independence* are widely used but little understood characterizations of cognitive style that are based on a person's strategies of visual perception and that purportedly reflect the manner in which the individual attempts to solve problems. Field-independent people are considered to be more individualistic, more self-motivating, and less dependent on external cues

than field-dependent learners. For a review and critique of field dependence and independence, see Rose, "Narrowing the Mind." Such measures of cognitive style may be most reliable, as here, when they are used in conjunction with correlating measures of individual differences.

⁹While Foss examined important issues and reported interesting observations, the available report of her research is frustratingly incomplete in its methodological details and analysis of the results. The number of subjects in the study was small ($n=10$), and no inferential statistics were reported. It is therefore impossible to tell whether the differences described here were statistically significant.

¹⁰As the research described here suggests, a reader's (or writer's) tolerance for the confusion engendered by hypertext is clearly situation-specific. Readers' attitudes will vary on the basis of their reasons for reading, their cognitive abilities and familiarity with the content area, their deadlines, the type of texts involved, and so on. Many people enjoy some confusion when they read a modern novel—but not when they read a political treatise or an instruction manual. Hypertext, of course, has effects that go beyond thwarting comprehension on the most basic level. In chapter 9 of this volume, Moulthrop and Kaplan describe other consequences—some positive and some clearly negative—of the deliberate use of hypertext to break down the boundaries between writer and reader, between previously autonomous texts and between texts and interpretations. Their purpose was both playful and didactic—to encourage students to become active, "strong" readers of any text, even those presented in a seemingly closed, authoritative, linear, hardcover publication. As they discovered (and as Johnson-Eilola remarks in the overview of this section), hypertext may subvert even these intentions—silencing where it was intended to give voice.

¹¹This essay develops and extends ideas first presented at Hypertext '87 (Charney, "Comprehending Non-linear Text"). The research was generously supported by a Penn State Research Initiation Grant and by Tektronix, Inc. I am grateful to Christine Neuwirth, Rich Carlson, and Mark Detweiler for comments on earlier versions of the manuscript.