5

Using Hypertext Systems

OVERVIEW LAND CONTENTS

Chapters 5 and 6 address issues raised by the recent advent of computerized information systems and the growing popularity of the World Wide Web as a means to acquire documentary information for purposes of education, training, or personal development. Chapter 5 focuses on the concept of hypertext, which forms the backbone of current online information systems. Hypertext consists of linking pages of electronic text in a nonlinear way, by means of semantic links. The concept of hypertext was invented in the 1960s, but it was not until the advent of the World Wide Web that hypertext emerged as a leading paradigm in electronic publishing. Early theorists saw hypertext as a means to facilitate access to content information to a wide range of users, by letting users freely "navigate" as a function of their needs and purposes. Empirical studies have found that navigating hypertext is a complex cognitive activity. Furthermore, the effectiveness of hypertext systems varies as a function of linking structure and content representation. The design of explicit and unambiguous hypertext organizers (e.g., content maps) has emerged as a key issue for the usability of nonlinear information systems.

Contents

INTRODUCTION

Over the past 25 years, the advent of personal computing has initiated major transformations in many areas of human activity. In particular, computers have deeply transformed information-related activities in professional contexts through, for example, office computing, process control, or intranets. Computers have also made their way into the sphere of education, supporting both general purpose tasks (e.g., word-processing, spreadsheet computing) and specialized learning activities (e.g., computer-assisted learning, distance education, Web-based learning). Finally, computers have become more and more present in the private sphere, where people use them both for leisure activities and personal management. The advent of the Internet and the explosion of general-purpose online services have consecrated the advent of an "information society," at least among the wealthiest social groups from the more developed countries.

The surge of powerful and versatile computer technologies has stimulated enthusiastic beliefs and expectations as regards the capacity of those technologies to accelerate the sharing of knowledge, culture, and entertainment. Many also think that computers will help resolve issues like unemployment, education, community involvement, and safety in western postindustrial society. There is, however, very little scientific knowledge concerning the effects—in terms of cognition, but also social integration and well-being—of widespread computer networking and intensive computer use. This is because research on the social and psychological impacts of information technologies has been somewhat overshadowed by technological development and innovation, combined with a rather positive stance toward technology on the part of policymakers, the media, and the general public.

I do not aim to cover such a large and complex issue here. Instead, I concentrate on hypertext as the core organizing principle of computerized information technology. This and the following chapter both address the effects of hypertext on document comprehension and document use. This chapter concentrates on the basic cognitive processes at work in hypertext perusal, whereas chapter 6 takes a broader perspective on the uses of Web-based hypertext in educational contexts.

In the following sections, I start with a brief review of the concept of hypertext (for related works see Conklin, 1987; McKnight, Dillon, & Richardson, 1993; Nielsen, 1995; Rouet, Levonen, Dillon, & Spiro, 1996;

Van Oostendorp, 2003). Then I examine the mental processes at work when using hypertext for various purposes. Hypertext was designed in order to promote not just reading, but sophisticated interactions with large bodies of texts and documents. I study the potential and limitations of the new medium for comprehension and information search, and review the issue of prerequisite skills involved in hypertext use.

A growing body of studies brings converging evidence that the hypertext user needs a great deal of support in order to navigate with a manageable mental cost, to locate information quickly and efficiently, and to make sense of the proposed networking of pages. Contrary to what many believe, such a need for support may be even greater than with printed documents (see chapter 2). Therefore, in the last part of the chapter, I focus on the role of electronic content representation (e.g., menus and graphical overviews) and their effects on hypertext comprehension and use.

5.1. AN OVERVIEW OF RESEARCH INTO HYPERTEXT USES

Where Does Hypertext Come From? The history of hypertext is rooted in the explosion of the publishing industry, in the emergence of modern librarianship, as well as in utopian undertakings aimed at creating universal knowledge repositories throughout the 19th and 20th centuries, such as Paul Otlet's Mundaneum (Rayward, 1991). The principle of nonlinear associative linking of information, which is the hallmark of hypertext, is generally attributed to Bush's (1945) article "As We May Think." Bush reflected on a major concern in the precomputer contemporary times, that is, the tremendous increase in the amount of scientific information available to scientists, with the correlative problem of retrieving and consulting information relevant to one's purposes. Bush noted:

> The summation of human experience is being expanded at a prodigious rate, and the means we use for threading through the consequent maze to the momentarily important item is the same as was used in the days of square-rigged ships.

He further diagnosed:

> Our ineptitude in getting at the record is largely caused by the artificiality of systems of indexing. When data of any sort are placed in storage, they are filed alphabetically or numerically, and information is found (when it is) by tracing it down from subclass to subclass (...). Having found one item, moreover, one has to emerge from the system and re-enter on a new path.

Bush went on describing the "Memex," a device that would allow its user to create associative trails by which any two pieces of information could be tied to each other. Later on, the retrieval of one such piece of information would allow the user to consult all the elements of the trail. This device, Bush argued, would let scientists, librarians, engineers, and so forth, develop applications tailored to their information needs.

The term *hypertext* was coined some 20 years later by Ted Nelson, who considered the possibility of constructing worldwide electronic information networks for the sharing of knowledge (for a historical summary, see Nielsen, 1995). It was not until the end of the 1970s, however, that the first prototype hypertext systems were actually developed as computer software. The Notecards® Software, by the Xerox Parc, was among the earliest, followed by several others. Hypertext as a research and development area was consecrated in 1987 with the first international conference on hypertext in Chapel Hill (Smith, Halasz, Yankelovich, Schwartz, & Weiss, 1987). The first European Conference on Hypertext took place in Versailles in 1990 (Rizk, Streitz, & André, 1990). During the 1990s, the industry of hypertext grew at a very fast rate, boosted by the advent of multimedia personal computers and the explosion of the Internet.

What Is Hypertext? From a technical standpoint, hypertext is a computer database that contains textual information (Conklin, 1987). The information is organized in files, and each file is connected to one or several others by means of software links. Hypertext systems also include a user interface that allows the user to search, display, and navigate the database using input and output interfaces (i.e., keyboard, mouse, screen, and printer). Nowadays, a large majority of Web sites and Web-based services actually include hypertexts.

The advent of hypertext has represented a significant breakthrough in the area of document design. Hypertext allowed the creation of networked arrangements of information units, challenging the traditional codex, that is, the linear arrangement of pages in a printed volume. More important, the advent of hypertext allowed writers and publishers to dissociate the contents of a database from its actual display. It was, therefore, possible to propose several representations of the same materials, without affecting the materials themselves. This is to be compared to the publishing process in the printed world, where any change in the presentation of information requires the reconstruction of an entirely new printed object.

Perhaps the most emblematic feature of hypertext is the use of embedded menus, or hyperlinks. Until the mid-1980s, database interfacing was dominated by the command line and hierarchical menus, separated visually and functionally from database contents. With the popularization of graphical interfaces and point-and-click devices such as the mouse, systems were created in which content information and command (or navigation) information were mixed up. In one of the earliest papers on the topic, Koved and Shneiderman (1986) defined embedded menus as follows: "In embedded menus, highlighted words or phrases within the text become the menu items, and are selectable using the commonly used touch screen, cursor, and mouse methods (...)" (p. 312).

Twenty years later, any Internet user may check that the embedded hyperlink was certainly a good idea, as millions of Web pages do implement this type of navigation device.

Is There Anything Psychological to Hypertext? One may wonder, though, whether hypertext technologies deserve a specific section in a book concerned with the psychology of comprehension. After all, readers remain readers and the general rules of language processing should apply, whatever the medium. However, many hypertext researchers and designers have claimed that the unique features of hypertext can qualitatively affect the cognitive processes brought to bear when reading and comprehending. For example, hypertext presentation of multiple documents, with embedded links to related sources and search tools, may promote students' reasoning at the intertextual, not just textual level. Hypertext features may also enhance the selection of relevant passages, the comparison of information, and the establishment of relationships between different types of information. Research on hypertext design and hypertext use during the past 15 years has shown, however, that augmenting people's ability to interact meaningfully with texts and documents is a subtle art, and that hypertext is by no means a magical solution to the problem of information access and information comprehension.

Independent from the still-vivid debate over hypertext promises and pitfalls, hypertext is important from a psychological standpoint because it gives way to new types of content representation and navigation tools. Those tools may not have any direct comparison with those available on paper. They raise, however, new issues as regards the nature of text comprehension and information usage processes.

5.2. COGNITIVE ISSUES IN USING HYPERTEXT

In this section, I focus on the issue of how people read and understand information presented in the form of hypertexts. This issue is part of a broader research area that investigates the use of electronic texts and information systems (see Dillon, 1994; Marchionini, 1995). Research studies dealing specifically with hypertext systems may be grouped into two broad categories, depending on the type of activity involved. The first line of research has attempted to study the impact of hypertext presentation on user comprehension of the contents. The second line of research has examined the impact of hypertext on people's ability to retrieve information. These two lines of research must be studied separately because, as I tried to show in earlier chapters, text comprehension and information search rely on distinct cognitive processes. I summarize only a few studies within each line of research, inasmuch as more detailed reviews have already been published elsewhere (Chen & Rada, 1996; Dillon & Gabbard, 1998; Rouet, 1992; Rouet & Levonen, 1996). I focus on the cognitive skills and abilities that have been found to be related to user performance in hypertext-based activities. The broader issues of whether and how students may take advantage of hypertext in the context of learning activities is addressed in chapter 6.

5.2.1. Hypertext Reading and Hypertext Comprehension

Ever since hypertext became a concrete technology, researchers have been eager to demonstrate that the use of hypertext systems may have beneficial effects on readers' comprehension and learning. Some researchers have tried a rather conservative approach that consisted of using hyperlinks only to present adjunct information within a linear online text. Others have attempted to demonstrate the benefits of "true hypertexts," that is, hypertexts that contain networked pages of information.

Hyperlinks for Online Definitions. An obvious application of hypertext is the provision of online definitions to readers of lengthy and unfamiliar texts. Checking the meaning of a word in a printed text is often a bother because one does not always have a dictionary at hand. Even with a dictionary, searching the word, reading the definition, and returning to the text can be quite disrupting. Hypertext links may provide a convenient means to provide definitions quickly and at little cost. In one of the earliest controlled experiments involving hypertext, Lachman (1989) hypothesized that the online definition of unfamiliar words may improve the comprehension of an expository text presented on a computer display, but only if the defined words are important to the meaning of the text. Lachman asked 32 college students to study a 6,522-word psychology chapter presented on 28 screen-pages. On each page, the definition of one word or phrase could be called for. Based on Kintsch and van Dijk's (1978) theory, Lachman parsed the text into a macrostructure and a microstructure. For half the subjects, the selectable words were part of the text macrostructure (i.e., important words); for the other half, the selectable words were not part of the macrostructure (i.e., less important words).

Overall, the participants selected 70% of the available definitions. Furthermore, the proportion of selected items as well as the time taken to read the definitions were greater in the second half of the text, but only in the "important definitions" group. Finally, the "important definitions" group obtained higher scores on a comprehension posttest.

Thus, in order to be efficient, computerized assistance has to do more than just provide additional information: This information must trigger effective comprehension processes. University students calibrate their use of online assistance as a function of the assistance effectiveness. More generally, this study illustrates the importance of a psychological theory for the design of user-centered information systems. The macrostructure theory of Kintsch and van Dijk (1978; van Dijk, 1980) led to accurate predictions of which definitions are likely to improve the comprehension of a long expository text.

The way inserted definitions are displayed may also influence the readers' willingness to use them. Wright (1991) reported a series of experiments in which online definitions of unknown words were inserted in a computer-displayed text according to various display options. When definitions could be read directly by clicking highlighted words, 93% of the available definitions were selected. Not highlighting the definable words caused this proportion to drop to 61%. Finally, an intermediate proportion (76%) was obtained with definitions grouped in a separate "glossary" instead of being directly selectable (Black, Wright, Black, & Norman, 1992).

Overall, the evidence suggests that online, optional definitions are most efficient under three conditions: when subjects are mature readers (for younger readers, compulsory definitions seem to result in similar or better performance; see Reinking & Rickman, 1990); when the definitions concern terms that are important for the particular text considered; and when the defined terms are clearly signaled and immediately accessible, to avoid harmful disruptions of the comprehension processes.

Comprehension of Linear Versus Nonlinear Text. During the 1985–1995 decade, the production of hypertext and hypertext browsers increased by large numbers, and so did the number of published studies providing empirical comparisons of linear text and hypertext for comprehension (Chen & Rada, 1996). As pointed out by Dillon and Gabbard (1998), however, many such experiments lacked a rigorous design and/or a sufficient level of control to provide interpretable data. Furthermore, due to the very versatility of hypertext, it is difficult to compare results across studies because one has to take into account differences in the participants, tasks, materials, directions, dependent measures, and so forth. For this reason, the actual impacts of hypertext on reader comprehension is still open to debate.

One of the earliest empirical studies of hypertext versus linear text comprehension was conducted by Gordon, Gustavel, Moore, and Hankey (1988). They selected printed magazine and technical articles, and converted them into hypertext using "informal subjective judgment." In the hypertext version, a summary of the document was presented first. The readers could access complementary information by selecting keywords (i.e., hyperlinks). The authors tested the linear and hypertext versions in order to find out which one would yield the best comprehension performance. Twenty-four students with little prior experience of hypertext reading participated in the experiment. They had to perform two distinct reading tasks: read a magazine article as they would do for leisure, then read a technical article in order to understand its content. In the "leisure reading" situation, the linear version yielded a better recall of the text's important information. Moreover, the participants preferred the linear version, and found the hypertext version harder to use. In the more demanding reading task, however, the participants did not find any difference between the two versions. The authors concluded that hypertext was probably not suited for reading situations that do not specifically aim at comprehension or learning.

One particular problem evidenced in the Gordon et al. study was the participants' feeling of uncertainty when reading the hypertext presentation. Some participants mentioned that, when using the hypertext version, they were bothered because they did not know "what was behind the door," that is, what information they would get by selecting the keywords. This finding suggests that better information about the destination of hyperlinks (i.e., labeled links) may have facilitated hypertext navigation. But it may also be suggested that when readers do not have any specific reading purpose, they cannot really take advantage of the "freedom" to choose or not to choose information categories present in hypertext. Reading a linear text is then as pleasant and as rewarding as reading hypertext (Charney, 1994). On the other hand, when asked to read more intensely the technical paper, the participants may have been encouraged to visit the hyperlink contents more systematically.

Early studies of hypertext found that navigating nonlinear document structures presented new challenges to the reader. Foss (1989, experiment 1) noted that hypertext readers tended to "loop" in the hypertext, and to flip through pages instead of reading them carefully. Self-reports indicated that looping and flipping did not reflect deliberate strategies, but resulted from a disorientation problem. Subjects reported difficulties in defining an optimal reading order, and in locating themselves in the network (see also Edwards & Hardman, 1989). Another study by Foss (1989, experiment 2) also highlighted the problem of managing a nonlinear reading task. Foss asked 10 adults to use a geographical hypertext database in order to perform a task involving the display and comparison of several cards. Foss reported two main types of problems: First, some subjects made too few comparisons and tended to lose track of their hypotheses or to forget how they had come to a conclusion. This was interpreted as a "search strategy" problem, or not having a good representation of the task requirements. Second, some subjects opened too few or too many windows at the same time, and/or positioned the windows in a way that did not allow easy comparison. This was interpreted as a "task management" problem, or not knowing how to perform the task. The two problems are not independent. For instance, poor task management (e.g., opening too many windows at the same time on the computer desktop) may prevent subjects from applying a good task representation (e.g., reasoning by elimination). In other terms, a coherent representation of the environment (what information is available and how to access it) is essential for effective access to the information of interest (see also Egan, Remde, Landauer, Lochbaum, & Gomez, 1989).

The navigation problem was also apparent in hypertext users' comments. Gray (1990) asked 10 students to read a 68-unit hypertext with the goal of answering questions. Think-aloud protocols were recorded during hypertext navigation and matched to the subjects' selections in the hypertext. The participants experienced several types of navigation problems: Some could not remember what they had read and not read, they missed organizational cues normally present in lengthy text, and they were not sure where to find the information they needed. When asked to draw a representation of the hypertext structure, subjects tended to reproduce conventional patterns: Sequences, simple hierarchies, or tables, rather than the actual hypertext layout. Gray concluded that novice hypertext users need analogies with conventional structures. She also suggested that with some training, hypertext users might become able to deal with loosely structured materials.

Hypertext navigation problems were observed even in very simple hypertexts. In my doctoral research work (Rouet, 1990, 1991), I asked French middle school students to read a hypertext made of six text passages connected to a single menu. Students were instructed to browse the hypertext until they had visited each unit at least once. Students' navigation patterns varied in the number and order of text selections. Some students read each unit just once, in an order that reflected semantic relations between topics. Other students went back several times to the same units ("looping"), and did not follow the relations between units ("jumping"). Results indicated that looping and jumping did not result from deliberate strategies, but reflected students' disorientation. Looping decreased when navigation was made easier by marking previous selections or making the relations between units explicit. Furthermore, in a second session, the students' selections followed more closely the relations across units, which suggests that familiarity did influence their navigation strategies.

These problems do not seem to be attributable to limitations or defects of early hypertext systems, as more recent studies have pointed out similar phenomena. For instance, Lee and Tedder (2003) found detrimental effects of a hypertext presentation of a history text, compared to a linear presentation, on students' factual comprehension of the contents. They suggested that hypertext reading creates a higher cognitive load on the reader and hence reduces the reader's ability to memorize content information. The reader must remember his or her location in the network, make decisions about where to go next, and keep track of pages previously visited (Wright, 1991). Given such constraints, it is hardly surprising that empirical comparisons between paper presentation (a familiar situation) and hypertext (a new, cognitively demanding situation) did not always favor hypertext (Dillon & Gabbard, 1998).

Hypertext and Relational Processing. How can we interpret the apparent "cognitive overload" that appears when reading and comprehending hypertext? Wenger and Payne (1996) suggested that the problem has to do with readers' comprehension strategies. To them, reading nonlinear materials fosters deeper relational processing on the part of the reader because the reader has to make a navigation decision after each page. The deeper relational processing causes an extra load on the reader's working memory, but it could be beneficial for texts that do not normally foster this type of processing. Hypertext presentation of descriptive materials could then force readers to establish connections between text units.

Wenger and Payne conducted two experiments in order to study the impact of hypertext on relational processing. In the first experiment, they used a secondary task technique in order to assess item-specific versus relational processing in hypertext. They compared the effects of two distinct secondary tasks. In the verbal-numeric task, the subject was asked to retain a series of six digits while reading each page of the text; in the spatial task, the subject had to hold in mind a configuration of six points on a four-by-four matrix. The latter task is assumed to require deeper relational processing, and thus it should draw more heavily on the subject's resources for relational processing. Twelve texts were prepared, based on scientific and technical publications. Six had a causal structure; six had a descriptive structure (following Meyer's 1985 typology; see chapter 2). The texts were divided into information "nodes" based on a thematic content analysis. The authors created semantic links between nodes. In the linear version, a unique link was drawn between each node and the next one, according to the basic rhetorical structure of the text. In the hypertext version, each node was linked to its superordinate theme, to its neighbors within the same global theme, and to other nodes in the text whenever this was justified. The presentation of the linear text started with the first passage in the source text; the presentation of the hypertext started with an alphabetic index where subjects could return during reading.

Forty university students read either the six causal texts or the six descriptive texts. Each subject read one text in each condition of presentation (linear, hypertext) and secondary task (numeric, spatial, or control). In the control condition, the series of digits was presented immediately after reading the page, just before the recall test. The analysis of subjects' free recall showed a better performance in the hypertext condition. Moreover, there was an interaction between presentation format, text type, and secondary task. Hypertext was better than linear text for descriptions, under a numeric secondary task condition. As regards the comprehension task, hypertext presentation was better for descriptive texts, whereas linear presentation was better for causal texts. Experiment 2 replicated experiment 1, using texts whose content was more familiar to the participants. Again, hypertext increased free recall of descriptions under the numeric secondary task.

Wenger and Payne's (1996) studies suggest that hypertext does not require more resources than linear text, but rather a qualitatively different type of resources. Hypertext may require more relational resources, which draw on Baddeley's (1986) "visuospatial sketchpad" in working memory. It should be noted, however, that these conclusions rest in large part on the analysis of free recall, which may not capture all the levels of processing involved in comprehension. Thus, hypertext reading may have simply encouraged the subjects to pay more attention to the materials, resulting in a better recall. Moreover, the experiments used a clearly structured text and an explicit linking scheme, mostly a hierarchical one. This may have greatly reduced the cognitive load and disorientation effects observed in other studies. In fact, other studies have found that hypertext linking with a predominant hierarchical structure does not harm comprehension, compared to linear chaining of the pages (Calisir & Gurel, 2003).

In conclusion, reading hypertext may stimulate the use of cognitive resources that are not used spontaneously when reading linear text, hence, maybe, a sense of greater effort, but also deeper processing of the materials. Thus, there is no simple linear relation between the amount of mental effort invested in the activity and the outcome in terms of comprehension. Even though there is no direct evidence for this, the data suggest an inverted U-shaped relationship, where either the lack of effort or an excessive level of effort both lead to nonoptimal performances. In all cases, the linking structure of the materials must reflect the actual semantic structure of the text. It should also be visible and readily interpretable by readers.

5.2.2. Searching for Information in Hypertext

In chapter 4, I argued that searching for information in documents is a complex activity that requires cognitive processes and strategies partly distinct from those involved in continuous reading. The construction of a mental model of the search task, and the management of the task requirements in working memory while searching, play a critical part in search effectiveness, just as does the efficient evaluation and selection of information categories, or planning (Dreher & Guthrie, 1990; Rouet, 2003). But search success is also tightly related to the quality of documents and content representation devices. Because hypertext came with new content representation and information search devices, many thought that hypertext may facilitate information search, compared to traditional, printed documents. In fact, hypertext was primarily designed to facilitate readers' interactions with large textual databases, that is, their selective access to relevant information.

Weyer's (1982) "dynamic book" may have been one of the earliest attempts to assess the effectiveness of hypertext for information search (Fig. 5.1). The principle of a dynamic book was to preserve the linear nature of the written text, while providing the reader with a set of tools to facilitate the retrieval and selection of related information. Weyer designed a prototype application based on a published high school history textbook. In addition to the contents of the textbook, Weyer's "Dynabook" included a variety of content representation devices and various search tools. The interface of the Dynabook included four major

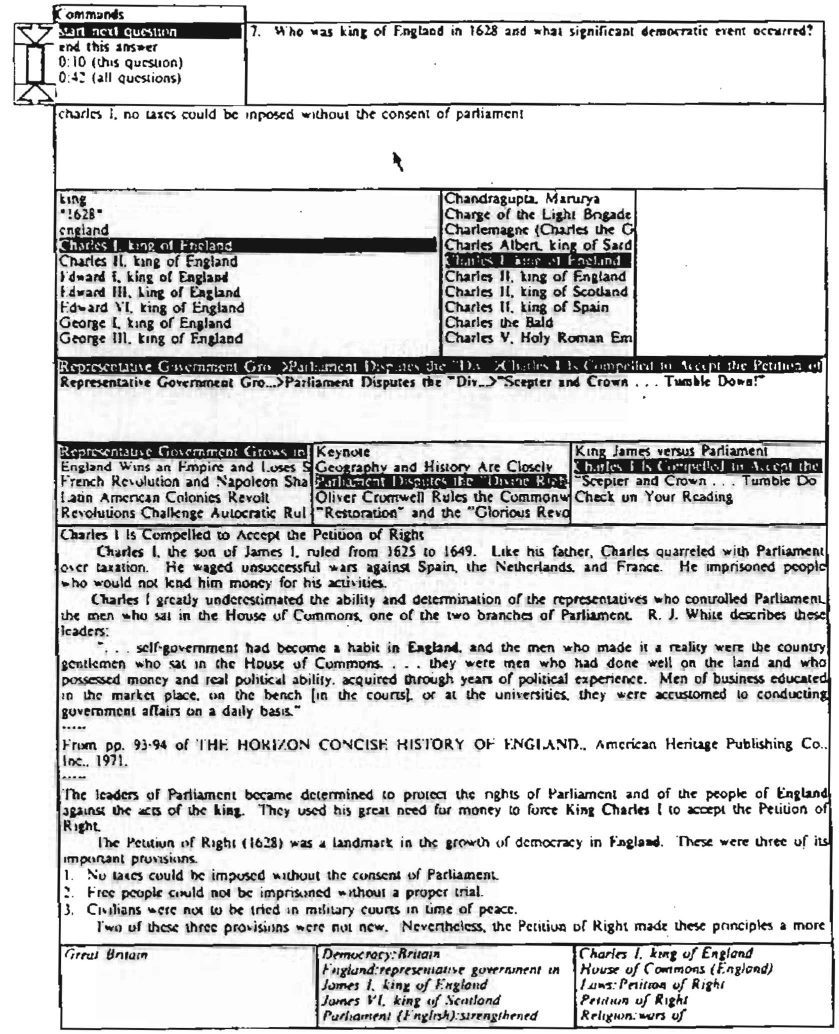

FIG. 5.1. The Dynamic Book (from Weyer, 1982). Reprinted by permission of Elsevier.

areas: the *command area* allowed users to read questions, keep track of time, and type in answers; the *subject area* displayed lists of subjects and subsubjects, and allowed users to type in their own keywords; the *title area* displayed topics and subtopics corresponding to the section of the manual currently displayed; and the *text area* displayed content information from the book. In addition, the prototype included two special navigation features: a hierarchical table of contents that enabled direct access to chapters, sections, and subsections, and a string pattern matching facility to enable easy information search.

In addition to introducing a very innovative prototype at the time, Weyer's study included an empirical test aimed at checking the usability and instructional value of the new device. To test the efficacy of the dynamic book, Weyer asked sixteen 8th-grade students (14-year-olds) to answer two series of 20 questions about the history of England. The students worked in pairs during two 2-hour sessions. In the first session, the students used a printout of the textbook in order to answer the first series of 20 questions. In the second session, the students used either a simple page-turning electronic version of the book or the full dynamic book in order to answer the second series of questions. Weyer reported qualitative observations showing that the students had quite some trouble managing the search task with the dynamic book. For instance, the students tended to mix up the different windows, and they could not make sense of some advanced features such as the addition of items to a table of cross-references. Furthermore, keyword search proved to be a difficult task, as students faced spelling problems and could hardly generate new keywords to describe a search objective. For instance, when searching information about the "French government," students insisted on finding the word "French" in the index, but they did not think of trying "France" or even "government." On other occasions they omitted important keywords, for instance, the keyword "holy" in a question about the "Holy Roman Empire," which led them to irrelevant sections of the manual (i.e., in this case, the Roman Empire).

When using the dynamic book, students also demonstrated problems due to their lack of knowledge of some high-level content representation devices. They tended to confound navigation in the text and navigation in the index, which could result in unexpected trouble. For instance, when trying to find out "who was king of England in 1628," a student performed a pattern-matching search using the phrase "king of England." As a result, the system displayed a list of kings that the student selected one after the other. When looking up sections about kings that ruled shortly before or after the target period, the student could have looked a bit forward or backward in the book, but he did not do so. Instead, the student continued to examine the list of kings until he found the correct answer (Charles I). Weyer also noted that searching the dynamic book to answer complex questions (e.g., questions that involved comparing information at various locations in the text) remained a complex cognitive task for high school students. The students seldom used high-level planning tools, such as the en route marking up of potentially interesting sections, for purposes of further reference.

The difficulties faced by participants may be interpreted in terms of the prerequisite knowledge and skills for document search (see chapter 4). First, middle school students are not fully aware of the role of textual organizers (e.g., tables, index), which form the backbone of the dynamic book. Next, the generation of appropriate search phrases requires a rich and flexible vocabulary, which typically students do not possess. Finally, information search imposes a heavy load in terms of memory management. Weyer's pioneering work showed that innovative information technology does not automatically eradicate those problems. In other words, although the dynamic book offered sophisticated search tools, it did not reduce the cognitive complexity of the search task. Weyer concluded that novice users might need some training before they become able to use the advanced features of the dynamic book efficiently.

A few years later, Shneiderman (1987) and his colleagues developed the Hyperties system, a hypertext system that was used for several applications, including a database on European history. In the Hyperties system, content information was categorized into a set of topics. Each topic was presented as a passage of text. Within each passage, several keywords were highlighted and could be selected, which led the reader to related topics in the database (see Koved & Shneiderman, 1986). A version of the database containing 106 articles was compared to its paper equivalent for question answering tasks. The paper version resulted in faster search for simple fact-retrieval questions. However, the users of the hypertext version performed equally rapidly for more complex items (Marchionini & Shneiderman, 1988; Shneiderman, 1989).

The Hyperties system was used in several other empirical studies. For instance, Wang and Liebscher (1988) asked university students to perform a series of search tasks using a version of the Hyperties system that included both embedded hyperlinks and an alphabetic index. The first experiment showed no difference in search effectiveness between the hypertext and a paper version of the document database. In their second experiment, Wang and Liebscher compared search using the index and search using the embedded hyperlinks. They found no overall difference. In the hyperlinks condition, however, search time decreased over trials, which suggests that the participants needed to get more familiar with this particular way of searching through a document database.

The lack of familiarity with hypertext search generally had negative consequences on users' performance. Gray and Shasha (1989) evaluated a database presenting information about sociology. They asked 60 university students to perform a series of five search tasks using one of three versions of the database: a version with a structured search facility, a version using the structured search facility plus a system of embedded links with categorized labels (e.g., "example," "comparison"), and a printed version with neither structured search nor embedded links. They found that the printed version allowed faster search for simple, explicit questions. For more complex questions, there was no difference across versions. They also noted that in the two computer conditions, search time tended to decrease across trials, suggesting again that search effectiveness could improve with the participants' familiarity with the system. McKnight, Dillon, and Richardson (1990) also reported evidence for disorientation problems in a hypertext search task. They asked a group of 16 adults to answer a series of 12 questions by searching a 40-card document. The document was presented in one of four formats: two hypertexts and two linear formats (paper and word processor). Search time was similar in the four conditions, but linear documents resulted in better answers. In the hypertext conditions, the subjects spent a greater proportion of time searching the menus, and they rarely used the direct links between cards. The authors concluded that inexperienced hypertext users face a task management problem, not knowing when and how to use the new navigation facilities.

To summarize the findings of both comprehension and information search experiments: Hypertext and related technologies have brought new means to represent and navigate complex information. The new technical tools did not, however, automatically result in observable benefits for the lay user. Instead, new problems were found, such as disorientation or poor task management. Very often, users seemed to lack the prerequisite skills or knowledge needed to take advantage of the system. This has resulted in an increased attention to the nature of the cognitive skills involved in using hypertexts and the factors underlying individual differences.

5.2.3. Prerequisite Cognitive Skills in Hypertext Use

Experiments on hypertext usage have often found large differences across individuals in terms of task management and performance. Such differences are attributable to a number of factors that have not been fully disentangled in the literature. Norman (1991) proposed a distinction between "inherent differences" (e.g., perceptual ability, memory capacity) and "acquired differences" (e.g., specific knowledge of the subject matter domain). Marchionini (1995) defined four factors of a person's ability to interact with computerized information systems: cognitive skills, domain knowledge, system knowledge, and searching knowledge. Empirical studies have confirmed the impact of both general cognitive dimensions and acquired knowledge and skills on people's use of hypertext (Marchionini, Dwiggins, Katz, & Lin, 1993).

Just like any complex mental activity, hypertext use depends on the efficiency of people's perceptual and cognitive processes. Particularly important are the processes that allow one to structure information received from the visual environment. Kim and Hirtle (1995) discussed the analogy between hypertext perusal and navigation in physical spaces. They argued that research on navigation in physical environments can be used as a reference frame to interpret the problems observed with hypertext navigation. Regardless of the specific task at hand, using a hypertext involves several cognitive activities, such as planning and executing a route through the network, processing content information, and coordinating the first two activities. Planning and executing routes in a hypertext network is in large part analogous to moving through a physical space. In both cases, routes may be based on one's knowledge of landmarks, on one's knowledge of familiar itineraries, or on one's learned map of the environment (survey-type of knowledge).

People vary in their ability to mentally construct and manipulate spatial representations, a skill called *spatial visualization*. Many studies have found an impact of spatial visualization on hypertext use (Chen & Czerwinsky, 1997; Lin, 2003; Westerman, Davies, Glendon, Stammers, & Matthews, 1995). Downing, Moore, and Brown (2005) found evidence for a relationship between spatial visualization and subjects' performance in a rather naturalistic bibliographical search task using the First Search online bibliographic tool. The time needed to access the first relevant article was shorter in participants with a high spatial visualization ability, independent from their expertise in the search domain. Freudenthal (2001) also found that a spatial ability measure predicted selection latencies at the deeper levels of a hierarchical menu. This relationship may explain older adults' lesser performances when using deep, as opposed to shallow, menu structures (see section 5.3.2).

Some researchers, however, have challenged the view that hypertext navigation relies on the same mental processes as navigation in physical spaces. Dillon, McKnight, and Richardson (1993) argued that the "space" metaphor is limited because hypertext is primarily structured according to semantic, not spatial dimensions. Farris, Jones, and Elgin (2002) argued that "hypermedia is inherently non-spatial" (p. 489), because it does not possess the qualities of depth and direction. They pointed out that there is no actual movement when using a hypermedia system. Thus, the exploration of hypermedia cannot result in the perception of spatial information. They deemed it unlikely that hypermedia users build a mental representation of the relative locations and depth of the pages that make up the hypermedia network. They designed several versions of a hypertext database of computer graphics that varied in depth (or the number of selection levels to be taken in order to reach a particular page) while sharing the same categorical arrangement. Forty university students explored the Web site for 5 minutes. Then they had to draw the structure of pages and links. The participants' drawings tended to reflect the categorical organization of the pictures rather than the connection structure of the hypertext. Farris et al.'s findings should be interpreted with caution since they are based on quite specific materials and on a task that is sensitive to people's generic schemata about information organization (see, e.g., Gray, 1990). They show, however, that it is not yet clear how spatial processes interact with other, semantic processes when exploring and using a hypertext system.

Be that as it may, the relationship between people's visuospatial ability and their performance in hypertext-based tasks is theoretically supported by text comprehension research. People tend to perform mental simulations of the "scenes" described in texts (see chapter 1). Furthermore, the use of graphical analogies helps people structure materials, even when the text does not have a strong spatial component (e.g., Glenberg & Langston, 1992). Robinson, Robinson, and Katayama (1999) demonstrated that the mental processing of graphical representations relies on the visuospatial component of working memory. Thus, spatial visualization skills might be related to hypertext use just as they are related to text comprehension in general (Graff, 2005).

Other cognitive dimensions also affect people's ability to use hypertext. Cognitive style, and especially field dependence versus independence, affects people's navigational style and efficiency. Field independence represents' people ability to reason independently from salient features in the perceptual environment. Kim (2001) found that field-independent students developed more efficient navigation strategies as they searched a Web site for specific information. However, the field dependence versus independence dimension was mostly significant for novice Web users (see also Chen & Macredie, 2004). Gillingham (1993) found a relationship between reading ability and hypertext search strategies. The more successful readers chose important hypertext nodes more often and read them relatively longer than unsuccessful readers. In addition to cognitive variables, the availability of prior domain knowledge, experience with search tasks in general, and experience with the particular search environment also improve people's use of hypertext.

Two major findings emerge from hypertext usage studies. First, the advantages of searching for information in hypertext are not immediately apparent in studies involving novice or inexperienced users. It was often noted that participants did not immediately understand how they could best take advantage of the system features. A training phase was necessary in order to assess the quality of the system properly. Second, some of the most advanced search or linking features available in hypertext may not be needed or even desirable. As Weyer (1982) noted, "Not having a feature may be better than using it badly" (p. 101). In fact, it seems that many of the prerequisites of effective search in traditional documents, such as knowing about metatextual organizers and how to use them effectively, are also present in search tasks involving electronic information systems.

5.3. REPRESENTING THE SEMANTIC STRUCTURE IN HYPERTEXT

From the early studies on, the history of hypertext systems has been that of a quest for *cognitive compatibility*, that is, a good level of match between the technical features of the system and the skills and needs of the users (Streitz, 1987). Cognitive compatibility can be achieved by selecting and representing relational information carefully (through hyperlinks), and by providing the user with structured top-level representations of the hypertext contents, or so-called *content maps*.

5.3.1. Link Labeling and Link Organization

Unlike printed documents, the top-level structure in hypertext is not made explicit through the linear arrangement of pages or chapters. Instead, structure is conveyed through semantic links that connect the hypertext nodes together. Hyperlinks allow the hypertext reader to move on from one passage to another. Critical issues in hypertext design are to find out what links should be offered with each hypertext node, how links should be labelled, and where they should be presented in the hypertext system.

What links should be included in a hypertext? Any content word or phrase on a hypertext page is a candidate for linking, provided that there exists another page dealing with contents related to that phrase. The number of potentially "linkable" items in a hypertext depends on a number of factors such as the amount of information contained in the hypertext, the way the hypertext is broken down into pages, and whether or not it is connected to other hypertexts. With the advent of Web-based hypertext, virtually any word or expression may be linked to either another page within the same hypertext (internal link) or to a page in another, external hypertext. Even though the concept of hyperlink is extremely seductive, the outcomes do not always live up to one's hopes.

Carelessly converting a linear text into a network of hypertext "cards" may decrease comprehensibility of the materials, especially due to coherence breakdowns at the local and global levels. Foltz (1996) pointed out that writers of traditional documents usually ensure local coherence by making contiguous sentences and passages share common referents. At the global level, texts and complex documents are usually arranged so that the reader can identify the overarching idea or theme (see also chapter 2). Foltz (1996) further pointed out that even though many hypertext designers have been aware of the need for local and global coherence, their approaches were rather empirical and a-theoretical. Foltz analyzed the paths followed by university students when exploring a 6,000-word hypertext for knowledge acquisition in introductory economics. The hypertext included both hierarchical and cross-section links. Using propositional analysis (see chapter 1), Foltz was able to identify those transitions that maintained textual coherence, and those that did not. Looking at how students with little initial knowledge explored the hypertext, Foltz reported that 80% to 90% of the transitions were respectful of the text's macrostructure. That is, the participants seldom used links that would lead them to remote parts of the hypertext. Foltz also observed that the more coherent the student's route was, the greater the amount of information the student was able to recall. In the second experiment, Foltz (1996) used verbal protocols to confirm that while browsing the hypertext, the students were busy keeping up with the text's macrostructure. He concluded that in order to understand a complex text, readers use a problem-solving approach that consists of building a route that will cause minimal disruption in their global representation of the text's content. Consequently, when linking information pages in hypertext systems, the designers should make sure that the subsequent action of jumping from the source page to the linked page will not cause excessive coherence breaks. This aspect of hypertext is especially crucial for information systems aimed at readers who are not experts in the domain.

Defining the global relations that may exist among portions of large documents is not an easy task, however. Holt and Howell (1992) called for rational methods to express semantic relationships between hypertext nodes. They designed a prototype hypertext generator, HyperNet, that allowed the authors to label all the links created among text nodes. In pilot tests, they found that authors did not have trouble linking text passages, but did have more trouble making explicit the semantic connections across passages. Therefore, they tried to gather a corpus of representative link types by asking university students to name the relationship between consecutive or unrelated pairs of paragraphs taken from a computer science manual. The students recognized a relationship in 79% of the consecutive pairs versus 19% of the unrelated pairs. Unfortunately, the students used rather vague and general expressions to characterize the relationships, for example, "how it works," "representation of," "example of," "description of." Thus, the provision of clear and explicit linking seems to require a great deal of expertise of the hypertext contents, in addition to careful consideration of the readers' needs.

Embedded Versus Explicit Menus. Embedded menus, or the integration of hyperlinks within the informational content of a hypertext, is one of the hallmarks of hypertexts. In their early discussion of the concept, Koved and Shneiderman (1986) argued that embedded menus may enhance navigation in computerized databases because they save screen space, and they preserve the semantic context in which a key word or phrase appears. They listed several applications of embedded menus in on-line databases, catalogs, spelling checkers, and programming editors. They briefly reviewed experimental results showing mixed but rather positive evidence in favor of embedded menus.

Subsequent empirical studies, however, failed to provide strong support in favor of embeddedness. Bernard, Hull, and Drake (2001) examined whether the location of links on a Web page had an influence on readers' performance at locating information, and on their subjective evaluation of document quality. They designed four 2-level hypertexts borrowed from *Scientific American* online. Each of the hypertexts used a different presentation of links. In version 1, the links were embedded within page 1 of the document. In version 2, explicit links were put at the bottom of the page. Version 3 placed explicit links at the top-left of the document, whereas version 4 placed the explicit links in the left margin, at the height of the corresponding content information (Fig. 5.2). Twenty volunteer students performed 10 search tasks with each version (content and presentation order were counterbalanced). There was no difference in search accuracy, time, or economy across versions. Embedded links received higher ratings in ease of navigation and ability to recognize key information, while bottom links had lowest ratings in comprehensibility and ability to follow main idea. Overall, embedded links were preferred most often, whereas bottom links were never preferred.

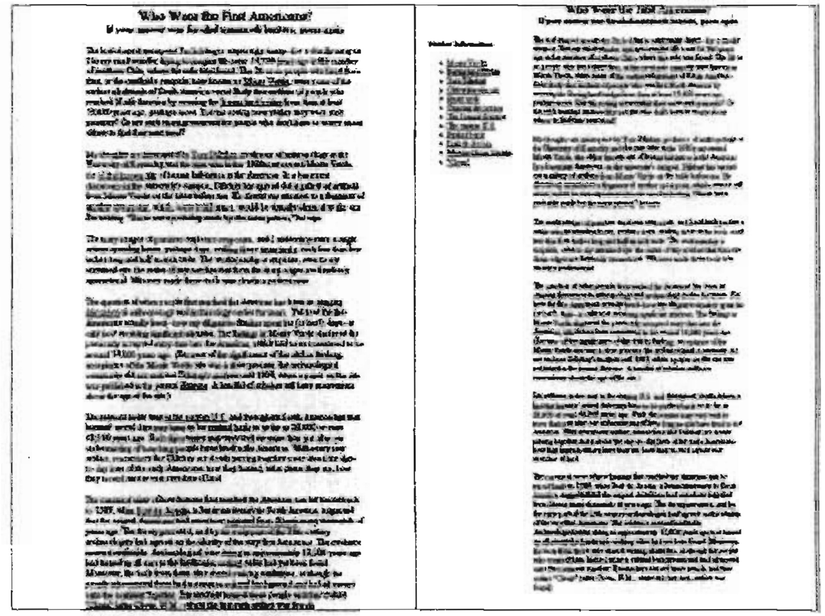

FIG. 5.2. Two versions of the hypertext used in the study by Bernard, Hull, and Drake (2001). To the left, a page with embedded links ("version 1"); to the right, same page with an explicit menu ("version 3"). Reprinted with the first author's kind permission.

Link Density. Link density, or the optimal number of links that should be included in a page, is another important issue in hypertext design. Khan and Locatis (1998) examined the impact of link density and link display on high school students' retrieval of specific information in a Web-based hypertext system. The hypertext included the equivalent of 15 double-spaced pages organized in nine chapters, dealing with influence and suggestion. Four versions resulted from the combination of link density (low = 3 links per page, higher = 6 links per page), and link display (explicit links in the form of menus vs. embedded within paragraphs). Both the high and low density versions had links pointing to relevant subsections within each chapter, as well as irrelevant internal and extraneous links. Sixty-four high school students were assigned to one of the four versions. They had to locate the answers to six questions, whose wording was identical, similar, or different from the wording of links. List display facilitated all aspects of search performance, compared to paragraph (embedded) display. A lower density of links improved search accuracy (but not search speed) and promoted a strategy based on the ordering of tasks from simplest to most complex (implicit).

The authors conjectured that "when users have more choices, they also have more opportunities to make mistakes," whereas "paragraph display had negative effects probably because users were likely to read the text and/or had to extract the information in the links from the surrounding prose" (p. 180). They admitted, however, that paragraph display could be more beneficial for more difficult tasks by encouraging the user to connect links and content.

Hierarchical Versus Network Arrangements of Links. Some hypertext theorists have claimed that the networking of hypertext pages should favor a richer reading experience, by letting hypertext readers revisit the same pages from different perspectives. Empirical studies, however, have consistently found that the networking of pages was strongly related to the disorientation syndrome (Rouet & Levonen, 1996). In experiments where the linking structure was manipulated, evidence was generally in favor of simple, hierarchical structures. For instance, Mohageg (1992) manipulated the organization of a hypertext database containing information about six North African countries. The database was organized according to linear, hierarchical, network, or mixed formats. Sixty-four adult paid volunteers were assigned to one of the four versions. They used the database to answer questions involving two, four, or six nodes. The hierarchical version was searched faster than the network version. Mohageg noted that orientation was an obvious concern of most participants. For instance, they often preferred to backtrack along previously visited pages instead of using the "home" button to return to the cover page. Other, more recent studies concluded that a hierarchical organization was better, especially for senior users, due to the higher "spatial cognitive load" of network structures (Graff, 2005; Lin, 2003).

Thus, linking should preserve a coherent, explicit top-level organization. In addition, hypertext readers should be provided with content representations that help them identify the top level structure, their current position, and possible itineraries within the hypertext.

5.3.2. Content Representation in Hypertext

An important aspect of hypertext is the way contents are represented in the system. In printed documents, content representation is achieved through the use of tables of contents, indexes, and other signaling devices (see chapter 2). In electronic databases, contents are often represented in the form of hierarchical menus (Norman, 1991). The advent of the Internet has popularized the use of hierarchical menus in Web portals, providing access to hundreds of categories in potentially all areas of interest to the general public. Given a number of categories or pages to be represented in a menu system, several display strategies are available. The hierarchy may contain only a few categories at each level and several embedded levels of information. Alternatively, the hierarchy may contain more items per page and a smaller number of layers (Fig. 5.3). The former strategy results in "broad" menus, whereas the latter one results in "deep" menus.

Pages in broad menus are visually more complex, as in the main broad menu in Fig. 5.3 (left). On the other hand, broad menus allow the user to reach content pages more directly. Deep menu pages are visually simpler, but they force the user to make several selections in a row in order to reach the desired category (like, e.g., the submenu selection in the deep menu in Fig. 5.3, right).

In the current state of the art, menu design varies a lot from one hypertext to another. For example, some Web portals offer deep menus with only a few options available at each level, whereas others offer broad menus with many categories available at each level. In some cases, the items are listed alphabetically, whereas in other cases the items are grouped by semantic categories (e.g., finance, travel ...). Given the increasing importance of using the Web for information search activities, it is important to assess the effects of these design strategies on users' information-seeking performance.

Past research has found that selection in a menu can be facilitated if items are grouped according to semantic categories (Giroux, Bergeron, & Lamarche, 1987; Snowberry, Parkinson, & Sisson, 1983) and if the depth–breadth ratio is optimal (Kiger, 1984; Miller, 1981; Parkinson, Sisson, & Snowberry, 1985), with about eight items per selection level (see also Norman, 1991, chapter 8, for a review). Excessive breadth can cause visual clutter and prevent users from finding relevant items in the menu; whereas excessive depth causes a sense of disorientation and cognitive overhead.

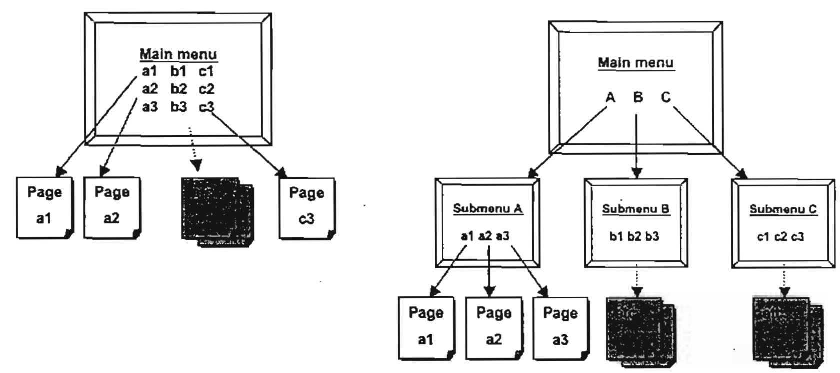

FIG. 5.3. Two types of menu organization: Broad menu (left), and Deep menu (right).

The need to compromise between breadth and depth appears to be independent from the technology, as studies conducted with recent, improved displays have essentially replicated the results of older studies. For example, Zaphiris, Shneiderman, & Norman (2002) compared traditional menus and "expandable" ones, in which subcategories appear in a pop-up area when the user selects a menu item. They found new evidence concerning the superiority of broad over deep menus. They also observed that search was generally faster with a traditional hierarchical menu than with an expandable menu, especially when the structure was deeper. The participants (21 university students) did not show any clear-cut preference for one type of menu. However, when the structure was deeper, they preferred the traditional menu over the expandable one. Zaphiris et al. concluded that expandable menus are acceptable only for shallow menu structures. Yu and Roh (2002) compared three types of menu presentation on university students' use and evaluation of a virtual shopping mall. They found that a pull-down menu that maximized visibility of intermediate categories was the most effective in terms of search speed for both specific and more global search tasks. Students' evaluation of design quality and disorientation, however, did not vary across menu types. Sears and Shneiderman (1994) pointed out the importance of ordering items within lists as a function of their selection frequency. "Split" menus, that present high-frequency items at the top, proved more effective in usability studies as well as a controlled experiment, compared to alphabetic menus.

Most of the empirical studies on menu search conducted so far have used students as participants. In contrast, there have not been so many studies of laypersons' use of menus in naturalistic activities. This was the purpose of a 2-year longitudinal study of Internet users' expectations, knowledge, and uses that my colleagues and I conducted in the urban area of Poitiers (France; see Rouet, 2005). As part of this study, Rouet, Ros, Jégou, and Metta (2003) examined the effects of menu design on younger and older adults' performance in a category search task involving various types of questions. Following previous studies, we hypothesized that deeper menus would decrease performance. We also expected that aging would negatively affect search performance, especially with deep menus and complex search probes (Freudenthal, 2001; Grahame, Laberge, & Scialfa, 2004; Westerman, 1995).

The participants were 50 volunteers from a panel of 100 laypersons participating in the longitudinal study. The sample included 9 men and 7 women aged 24–36 (younger adults); 5 men and 14 women aged 37–53 (intermediate adults); and 8 men and 7 women aged 54–80 (older adults). All the participants had been regular users of the Internet for over 18 months at the time of the experiment.

Rouet et al. (2003) designed a 400-item menu structure after existing Web portals. The menu structure presented a hierarchy of general interest categories and subcategories (e.g., education, travel, jobs, sports, and so forth). Three versions were developed: broad-categorized, broad-alphabetic, and deep. The broad menu versions involved a larger number of items per page (i.e., a maximum of 42), but only two levels of selection (main menu-submenus). The deep menu structure involved a maximum of only six items per page, but four levels of selection (Fig. 5.4).

In the categorized version of the broad menu, items were grouped according to semantic categories (e.g., "education," "travel," "jobs"; see Fig. 5.4a), whereas they were ordered alphabetically in the two other versions (Figs. 5.4b and 5.4c). The number and wording of target categories were identical across versions, as well as the visual characteristics of the display (e.g., size, color, etc.).

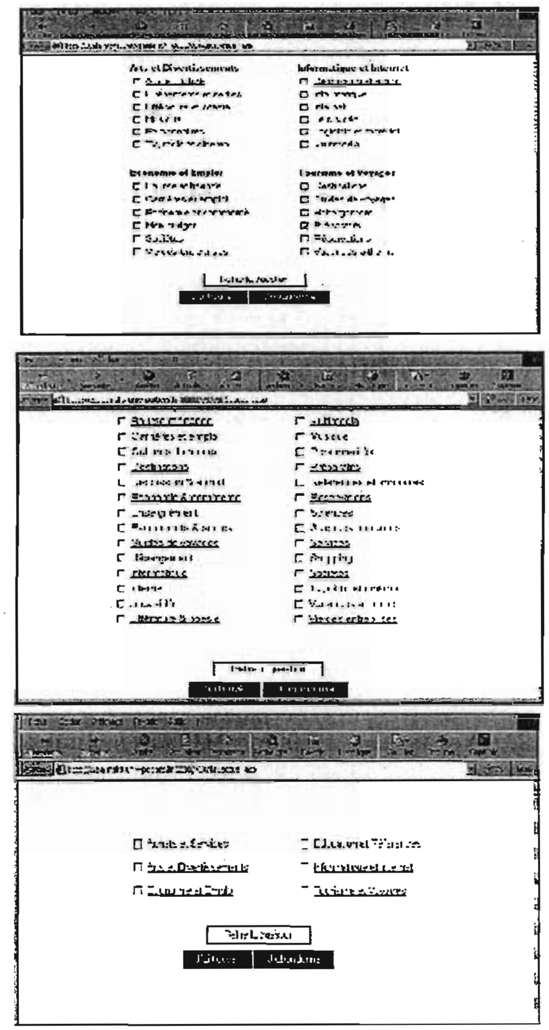

FIG. 5.4. The three menu organizations used in the study by Rouet, Ros, Jégou, and Metta (2003).

Twelve search questions were written based on keywords from the menu hierarchy, so as to be compatible with all three menu structures. For instance, the question "Find accommodation in a palace hotel" was written based on the hierarchy "tourism > accommodation > hotel > palace." Each question was written in four different versions in order to manipulate explicitness and length (Table 5.1). In the "implicit" versions, the original content words were replaced with synonyms (e.g., "Find lodging in a luxury inn"). In the "longer" versions, 2 or 3 content words were added to each question (e.g., "find temporary accommodation in a nice palace hotel in Paris"). Thus, for each question there was a short explicit, a short implicit, a long explicit, and a long implicit version.

All the participants were visited at their homes by appointment. They all possessed a standard PC with a 14-inch screen and a 56K modem connection to an Internet provider. The participants performed the 12 search tasks using each question presented in one of the four versions, and one of the three versions of the menu. For each question, the participants were asked to identify as accurately as they could the relevant subcategory using a check box. They could reread the question as many times as needed while searching. They were allowed to give up the search using a "give up" button. The categories selected, along with the selection delays, were automatically recorded and stored in the database by the Web server.

Despite the fact that all questions corresponded to a unique category, there was only a 52.5% average success rate. On 9% of the occasions, the participants selected another category in the correct submenu. On 38.5% of the occasions, they selected another category or failed to provide any answer. The broad categorized menu had a success rate slightly higher (57%) than both the deep and the broad alphabetic menus (50%). Both question explicitness and length had an impact on search success. Short explicit questions were answered in 67% of the cases, as opposed

TABLE 5.1
Example of Search Questions and Menu

Explicit, short:	
Find accommodation in a palace hotel	
Explicit, longer:	• Shopping
Find temporary accommodation in a nice palace hotel in Paris	• Arts and culture
	• Business and economy
	• Education
Implicit, short:	• Travel and accommodation
Find lodging in a luxury inn	• Computers and Internet
Implicit, longer:	
Find temporary lodging in a nice luxury inn in Paris	

Note. Based on Rouet, Ros, Jégou, & Metta (2003).

to 57%, 45%, and 41% for long explicit, short implicit, and long implicit questions, respectively. Finally, older participants tended to perform poorer when using the deep menus.

Rouet et al. also observed that participants reread the question less often with the broad-structured menu than in the other two conditions. Older adults also tended to reread fewer questions than the other groups, especially when searching the deep and broad-categorized menus. Looking back at the question may be interpreted as a measure of the cognitive load associated with searching a hierarchical document. Rouet (2003) found that the rate of question lookbacks increased when the question was complex or the domain was unfamiliar. The rate of rereading also increased with question length and complexity. The mental effort required to locate relevant categories seems influenced by the intrinsic complexity of the menu, as well as by the phrasing of the question.

A good way to help users retain a sense of orientation is to display the content representation permanently, at the left of the hypertext window. Jégou, Andréo, and Rouet (2001) trained 65 undergraduate psychology students to search a Web site in order to study course-related information. A third of the participants were novices in the use of Internet, whereas the remaining two thirds were occasional users. The Web site was designed so that two types of content representations were available: a table of contents and an index. Moreover, the content representation was always visible to the left of the screen, while content pages were displayed in the central area of the screen (Fig. 5.5).

The participants managed to locate relevant contents in more than 80% of the cases. The students reported positive feelings as regards their orientation and ease of navigation in the site, and they thought they would use this type of tool if it was made available as a learning resource. However, many students also reported that they had some trouble remembering the question while searching, and that they wished it was permanently displayed during the search. Thus, making available both the task representation (e.g., questions or study directions), a global representation of the information available, and the content information currently studied seems a condition for easy and efficient navigation in hypertext.

5.3.3. The Potential of Graphical Content Representations

Empirical studies emphasize the need to provide hypertext users with quality content representations. Because hypertext use resembles navigation in a physical space, many authors have recommended providing the user with metaphors (Hsu & Schwen, 2003) or content maps that represent the arrangement of nodes and links relevant to the user's purposes (Kim & Hirtle, 1995). Nilsson and Mayer (2002) defined content maps as:

> a graphic representation of a hypertext document, in which the pages of the document are represented by visual objects (whether simply the title of a page or an icon representing a page) and the links between pages are represented by lines or arrows connecting the visual objects. (p. 2)

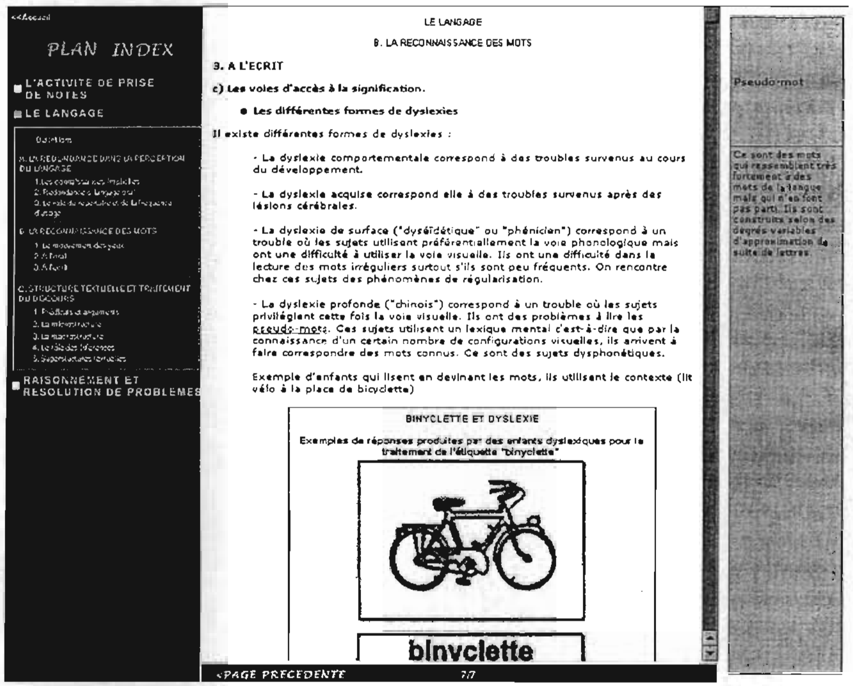

FIG. 5.5. A screenshot of the Web site used in the study by Jégou, Andréo, and Rouet (2001). The left side of the screen features the content representation (table of contents or index), the central part displays the current page, and the right-hand side area is for definitions or adjunct information that the user selects through hyperlinks.

In a concept map, each node represents the contents of a text passage by a thematic phrase. Links represent different kinds of relationships between concepts (Stanton, Taylor, & Tweedie, 1992). The reader accesses a hypertext section by selecting one of the nodes displayed in the concept map. Thus, concept maps serve two distinct purposes: They inform the reader about the contents of the hypertext, and they allow the reader to display information units on the computer screen.

Content Maps and Hypertext Comprehension. Empirical studies of the effects of interactive content representations on college students' comprehension of hypertext have had mixed results. Some studies found a beneficial effect of hierarchical content representations on hypertext comprehension. Dee-Lucas and Larkin (1995) hypothesized that hypertext presentation may facilitate the process of selectively reviewing a document because hypertext provides direct access to the document units. However, this would be the case only if the content map of the hypertext reflects the semantic structure of the domain. Dee-Lucas and Larkin compared three presentation formats of a nine-unit document on electricity: linear, unstructured hypertext (with an alphabetic index), and structured hypertext (with a hierarchical content map).

In their first experiment, 45 college students were asked to study the document in order to be prepared for a test on its content. The participants first read the document in a fixed order and then were allowed to review it. Compared to linear presentation, both hypertext formats resulted in more units being reviewed. However, the participants spent more time selecting units to be reviewed in the alphabetic index than in the hierarchical content map. There was no difference in the total amount of information recalled, but the two hypertext formats resulted in a larger "breadth" of recall: The subjects recalled more unit titles and ideas from more text units. Furthermore, the structured hypertext condition resulted in better memory for title locations in the index. In the second experiment, 63 college students read the same text in one of the same three conditions. However, they were given a specific reading objective—being able to summarize the document. The subjects reviewed more units than in experiment 1, and the differences between presentation formats were greatly reduced. The authors concluded that a hierarchical content map facilitates the construction of a text macrostructure, that is, a more integrated representation of the text content (see chapter 1). They also suggested that the demands of the task can override the effects of different presentation formats.

This study points out several interesting phenomena. First, hypertext presentation promoted the learning of the document structure, as evidenced by the larger "breadth" of recall in experiment 1. Happ and Stanners (1991) also reported that hypertext presentation led to a better learning of the structure of a relational conceptual system. Second, a hierarchical content map facilitated subjects' orientation in the hypertext. Students were faster at selecting the sections in the hierarchical than in the alphabetic index. This is consistent with Simpson and McKnight (1990), who observed facilitative effects of a hierarchical index compared to an alphabetic index. Subjects were also better at answering content questions and at reconstructing the hypertext structure. Third, the influence of different presentation formats varies according to task requirements. More demanding and/or more specific tasks may reduce the effects of content representation devices. For instance, in a study by McDonald and Stevenson (1996), undergraduate psychology students had to read a hypertext on the topic of human learning using one of three presentation formats (a linear text vs. a hierarchical map vs. a network map). Then, the students were asked to use the hypertext in order to answer 10 questions. The results failed to show any effect of the type of content representation on students' comprehension, as assessed by the number of questions correctly answered.

A potential drawback of rich external content representations is that they may decrease the user's effort to understand and memorize the hypertext's organization. Nilsson and Mayer (2002) discussed the potential effects of content maps in terms of two distinct theoretical constructs: the cognitive load theory and the active learning theory. According to the cognitive load theory, maps would help users locate relevant categories in a hypertext by providing an external representation of the structure, which saves the user the effort of building and retaining this representation in memory. According to the active learning view, however, using a map may reduce the amount of elaborative or constructive processes brought to bear by the learner while studying the document. This, in turn, may lead to a lesser learning of the hypertext contents.

Nilsson and Mayer designed a 150-page hypertext containing information about fish and other aquatic animals. The animals were categorized according to three different hierarchies: animal classification (e.g., rays and skates), area of habitat (e.g., Mediterranean open waters), and diet (e.g., omnivores). Three maps presenting each of the three hierarchies were included in the system. The maps could be accessed though explicit links displayed at the bottom of each animal page. The "No map" version did not include any such map.

In the first experiment, Nilsson and Mayer asked 53 undergraduate students to use the system either with or without maps in order to answer a series of 30 questions about aquatic animals (e.g., "How do ghost pipefish catch their food?"). The main finding was that even though participants in the map condition searched fewer pages during the initial trials, the number of pages searched in the no-map condition decreased more dramatically in subsequent trials. As a result, search in the no-map group was slightly more efficient in the last 10 trials. The authors also found that the participants in the no-map condition seemed to use a "task-based" strategy, that is, they considered the specific cues contained in each question, as opposed to a structural strategy based on their knowledge of the hypertext's overall organization. Nilsson and Mayer concluded that "any benefits from the structural information in the map was overshadowed by the negative effects of the map decreasing participants' involvement in the task" (p. 14). This experiment emphasized the importance of the reader's active involvement in a search task as a factor of success.

In the second experiment, Nilsson and Mayer replaced the content maps with a list of path-type expressions that showed the location of the current animal page, for example, "Habitats–Marine–Open waters–Mediterranean open waters–Common skate."

These expressions were "navigable," which means that the user could click on any of the component phrases (e.g., "open waters") to go directly to the corresponding group page. After a training period, path-type expressions speeded access to content pages, compared to a control condition. Spatial ability was positively related to search accuracy. Nilsson and Mayer concluded that the notion that any organizer would facilitate hypertext-based learning is inaccurate. They recommended that maps be used in systems aimed at one-time or occasional users, with no ambition of learning the system. The path-like organizer, on the other hand, seemed useful as a learning device, but only after a training period. Thus, depending on what the system is being designed for, different content representation and navigation tools are warranted.

Content Maps and Prior Knowledge. To further explain the effects of interactive content representations, other studies have considered readers' prior knowledge. Möller and Müller-Kalthoff (2000) showed that a hierarchical map facilitates comprehension only in low readers' prior knowledge. Undergraduate students with low prior versus high prior knowledge in psychology were asked to read a psychology text about cognitive models of writing, using one of two presentation formats (hierarchical content maps vs. no map) and to answer factual questions. The results showed a significant relationship between factual comprehension, hypertext organization, and prior knowledge. Only low prior knowledge readers gave better answers to questions with a hierarchical map.

Shapiro (1999) asked undergraduate students with low prior knowledge on ecosystems and high prior knowledge on animals' families to read a hypertext on these topics, with one of two learning goals: to learn specifically more about animal relationships in their ecosystem (goal A) or about similarities and differences between animals' families (goal B). The participants had to read the hypertext either using a hierarchical interactive map showing the categorical organization of species, or without the map. Then they had to answer explicit or implicit questions (i.e., questions about the textbase or on the situation model, respectively). The hierarchical map improved low prior knowledge participants' answers to implicit questions. The hierarchical map may have provided a conceptual aid to help low prior knowledge readers integrate new incoming information in their situation model. McDonald and Stevenson (1998) showed that navigation aids facilitate low prior knowledge (LK) readers' comprehension. Psychology students read a hypertext on the topic of discourse production. The text was presented either with a navigational aid (i.e., a network concept map or a simple list), or without such an aid (i.e., only as a set of hypertext nodes and links). Both types of aids lead to better comprehension, but only in LK students. Moreover, the time needed to answer questions was shorter when using a concept map than a list.

Another study, by Hofman and van Oostendorp (1999), tried to assess the effects of content representations on several representation levels as a function of readers' prior knowledge. Undergraduate students with high versus low prior knowledge were asked to study a science text on sun radiation and health. The text was presented either through a network concept map showing various types of relationships between ultraviolet radiation and skin cancer, or through an alphabetic topic list (i.e., without explicit high-level relations). Structural levels of text information (i.e., microstructure vs. macrostructure questions) as well as representation levels of text (i.e., textbase vs. situation model questions) were manipulated so as to produce four types of comprehension questions. Contrary to previous results, Hofman and van Oostendorp found that the concept map hindered LK students' situation model construction. They suggested that the concept map had diverted readers' attention from more appropriate levels of processing. For readers with little prior knowledge, simpler representations (e.g., content lists or hierarchical maps) may be more productive than complex network representations.

Thus, even though concept maps are usually thought to be beneficial, their actual effects vary across experiments. Potelle and Rouet (2003) suggested that the impact of content maps depends on an interaction between the features displayed in the map, on the one hand, and the user's prior knowledge of the domain, on the other hand. Hierarchical maps may facilitate the construction of the hypertext macrostructure in LK students (Dee-Lucas & Larkin, 1995; Shapiro, 1999) because they display basic global relationships among the topics dealt with in the text (Lorch & Lorch, 1995; see also chapter 2). Reading a hierarchical map may help LK students build a mental representation organized along categorical or thematic dimensions. On the other hand, a network map could hinder LK students' construction of the macrostructure because of its too complex semantic links (Hofman & van Oostendorp, 1999).

Potelle and Rouet (2003) designed a simple hypertext made of seven content cards about various aspects of "social influence," a core topic in social psychology studies. They designed three content representations of the hypertext (Fig. 5.6). The *hierarchical map* was organized with superordinate and subordinate links from the most general to the most specific topics about social influence. The *network map* was organized by connecting the main topics with semantic links. The relevant links were identified in a pilot study involving 10 PhD students who were asked to draw connections between two parallel lists of topics. Finally, the *alphabetic list* presented the topics in alphabetic order, without explicit connections.

Potelle and Rouet hypothesized that the hierarchical map would function as a structural cue for all the readers, which would improve comprehension, especially at a macrostructural level. The network map, however, was based on implicit semantic relations generated by expert students. Understanding these relations (e.g., "minority influence–innovation") requires some prior knowledge of the domain. Thus, having to study this type of overview might be detrimental to novice readers. Forty-seven students participated in the experiment. The participants were categorized as domain novices versus specialists based on the median split of a knowledge pretest. They studied the hypertext for a period of 20 minutes, with an explicit comprehension objective. Comprehension was assessed through a 16-item multiple choice questionnaire and a summary task. The participants were also asked to draw a map of the hypertext from memory.

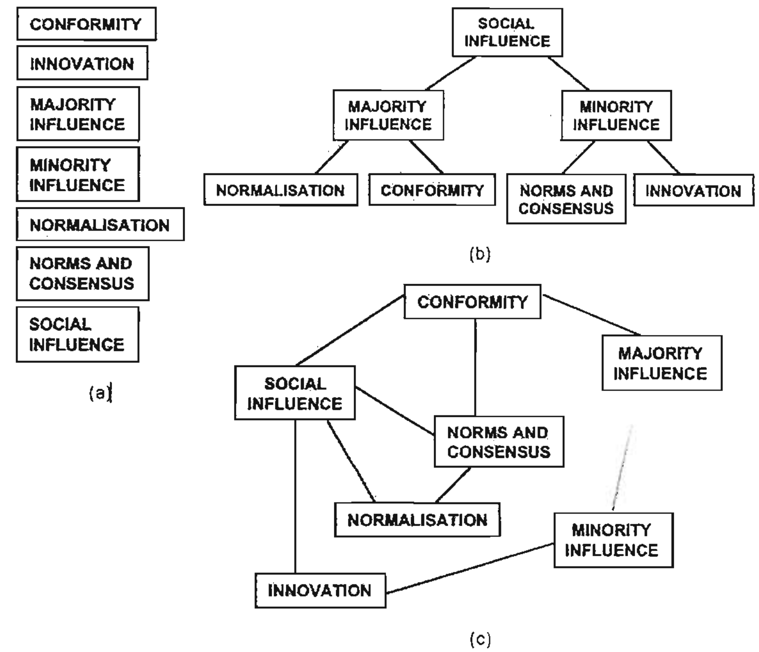

FIG. 5.6. Three content representations of the hypertext on "social influence" (based on the materials used by Potelle and Rouet, 2003); (a) alphabetic index; (b) hierarchical map; (c) semantic network.

Comprehension was positively related to prior knowledge, and microstructure questions were better answered than macrostructure questions. An interaction between prior knowledge level and the type of content representation was found. Low knowledge students had better scores when using the version of the hypertext with a hierarchical map than with the other two formats. For the more expert students, the type of content representation had no significant impact on comprehension. As expected, the effect was stronger for macrostructure than for microstructure questions, even though the three-way interaction failed to reach significance. Low knowledge students also included more thematic ideas in their summaries when reading from a hierarchical map. Finally, they drew more accurate maps of the hierarchical hypertext than with the other two systems. Again, no difference was found for high knowledge students.

Potelle and Rouet concluded that the effects of content representation depend in part on the reader's prior knowledge level. As for any text organizer (see chapter 2), the reader must be able to recognize and use the signals presented in a content map. If the map is ambiguous, or uses unfamiliar symbols, then the result will be an added burden on the reader, with dubious effects on comprehension and recall. When designing hypertext systems, care should be taken to adjust the level and type of structural information to the capabilities and needs of the user (see also Carmel, Crawford, & Chen, 1992).

In summary, graphical content representations are useful to the extent that they convey organization principles that the hypertext reader can readily integrate. Furthermore, content representations must be consistent with the hypertext's actual semantic structure. As pointed out by Dallal, Quibble, and Wyatt (2000), content representations, whether graphical or verbal, should help readers perceive local and global coherence links. This may be done through the use of various hypertextual organizers that play a complementary role: a global content representation, explicit links, indications as regards the context of a particular page, and so forth. Dallal et al. showed that the inclusion of such attributes improved students' ability to retrieve information in a Web site. They concluded that design guidelines based on cognitive research may lead to significant improvements in the quality of complex information systems.

CONCLUSIONS

Our ability to comprehend and use complex documents is partly dependent on the technologies that allow the production and dissemination of information. Since they were created, a few decades ago, hypertext technologies have had a deep impact in the area of document design and document publishing. Document writers and editors have gained much flexibility in the rules that govern the writing and displaying of large sets of texts. These new and exciting technical facilities have triggered optimistic expectations as to their potential role as facilitators of people's access to written information. However, so far hypertext has not always lived up to those expectations. A large number of experiments were conducted, and many have shown that novel ways of displaying and navigating information in hypertext had little or no impact on the readers' performance. Another form of evidence comes from the yet limited impact of hypertext in the publishing world. The explosion of multimedia technologies and the Internet did have a large impact on traditional publishing, especially in some specialized areas (e.g., science publishing). But it certainly did not cause any major drop in the production of traditional books, journals, magazines, and other printed materials. Furthermore, as the use of Web sites becomes more and more common in the general public, designers tend to rely on metaphors that are part of the users' background culture, instead of trying to force novel, artificial presentation formats into them (Nielsen & Tahir, 2002).

Powerful as it may be, hypertext is not always appropriate for presenting materials and tasks typical of the printed world. For instance, in many experiments the hypertext contents were adapted from printed documents, but the cues present in these documents were removed or replaced with less familiar ones. Moreover, initial hypertext studies

were conducted with novice hypertext users, who were probably tempted to look for the type of cues they would normally rely on. Literacy skills are acquired through years of education and practice, and it is likely that people need a bit of time before they can adjust to profound changes in their information environment. So far, the research suggests that online and printed information each have their areas of applications. Whether in print or online, what matters is document quality, and document quality may be obtained only by paying a great deal of consideration to why and how people use documents.

Meanwhile, both print and online text technologies are evolving rapidly. Despite the large body of theoretical and empirical studies published in the past decades, any final conclusion concerning the intrinsic benefits and limitations of hypertext would be very much premature. The technology is fast evolving and, more importantly, a new culture of hypertext usage is slowly emerging from the technological big bang that characterized the end of the 20th century. The advent of general-public Internet and World Wide Web services has given information designers many new opportunities to apply hypertext concepts. Many more questions have arisen, and many more theoretical and empirical studies remain to be conducted in order to find out about the potential of hypertext.

Recent research suggests that the road to usable hypertext goes through a deeper analysis of the rhetorical processes involved in nonlinear writing. Hypertext writers need to be more considerate and to anticipate potential comprehension problems in their readership. In particular, the content representation of complex documents must be designed carefully, so as to allow hypertext users to build up appropriate comprehension strategies. Once again, this traces back to the issue of complex texts and the prominent role of content representation and rhetorical cues (chapter 2). This issue is all the more important when considering the increasing use of open, Web-based hypertext systems in educational contexts.