CHAPTER

12

The Mental Web

Pedagogical and Cognitive Implications of the Net

JENNIFER WILEY
University of Illinois at Chicago

JONATHAN W. SCHOOLER
University of Pittsburgh

I. INTRODUCTION

Every new information technology from the first cave drawing to the invention of the alphabet and the advent of the television and the modern computer has sent ripples through society, forever altering ways that we communicate, interact, and perhaps even think. The emergence of the Internet (net) and the World Wide Web (web) is likely to be no exception, although the precise manner in which their impact is to be felt is still an area of marked speculation. In the domain of education, whether the emergence of the web and its uses in teaching and learning cause a true paradigm shift in educational practice, or a more modest reconceptualization, there are clearly major differences between teaching and learning from the web versus more traditional educational media. That, at least in principle, could have broad significance. In the following chapter, we briefly consider some of the potential pedagogical implications of differences between the web and both traditional classroom and text learning contexts. We then consider some areas that this review suggests are in particular need of further research. We close with a discussion of the pedagogical and cognitive implications of a

potentially fundamental aspect of the web, namely, its capacity to emulate and instantiate many of the processes of the mind.

II. DIFFERENCES BETWEEN THE WEB AND CLASSROOM LEARNING CONTEXTS

The present volume has enumerated many points of interest where traditional classroom practices diverge from the affordances of the web. The following is a summary of some of these central differences.

A. Physicality

The most obvious contrast between virtual media and traditional forms of education is the basic difference in physicality. In traditional classroom settings, the teacher and students are physically co-present, and texts are written in permanent form on pieces of paper. Some paper is compiled in books that are made available to students as assigned texts. Printed texts and conversation partners are physically present, and provide rich stimulation to the senses. However, the supply of books, partners, and perspectives is limited by physical space. In contrast, in virtual educational settings, discussion and information emerge from pixels on a computer screen. The physical context stays constant, a student interacts with a single computer and screen and keyboard, but the information that can be accessed is practically infinite. As a number of chapters have discussed, from these simple physical differences emerge many far-reaching implications.

B. Social Interactions

Perhaps the topic that has received the most attention has been the differences in social contexts between real and computerized settings. In the case of virtual learning environments, the lack of physical co-presence in itself alters the social context of the exchange of information. There has been great concern over the limited social interaction that may be available in virtual education. The HomeNet studies at Carnegie Mellon (Kraut et al. 1998) suggesting that people who spend too much time on-line tend to have disturbed interaction patterns, and are more likely to experience depression and loneliness, are often cited as evidence that computer use leads to negative consequences. But more recent studies, such as one reported by the Pew Internet and American Life Project (2000), have found different results. Women, in particular, who spend more time on-line tend to be better socially adjusted, as they often spend their time on-line connecting with distant relations and friends, and renewing old acquaintances. Even the

original HomeNet studies found that in some ways, and for some people, computer use supported social interaction in positive dimensions. Indeed as Reyna, Brainerd, Effken, Bootzin, and Lloyd (chapter 2) observe, there are good reasons to expect that as with many other aspects of the web, individual differences may play a critical role in determining the costs and benefits of the web relative to more traditional educational media. In particular, introverts, who can be intimidated by the social context of the classroom, may be able to contribute and even start friendships over the web in a manner that they might otherwise never have been able to achieve.

Other chapters also speak to the opportunities for on-line social interaction that may actually surpass live or "real" classroom interaction. For example, Sherman (chapter 6) traces the importance of social interaction in educational contexts in a review of literature on cooperative learning. He identifies several on-line tools that can be recruited to support cooperation and collaboration in computerized learning environments. He points out that these on-line tools can substitute for "real" interaction in a number of ways, and offers evidence that they are being used with increasing frequency, especially for distance learning courses.

Riva (chapter 7) takes the notion of on-line communication one step further by examining the differences between on-line communication and virtual discussions, and exploring notions of how the unique capabilities of on-line interactions may be best understood. Instead of merely viewing communication, and a sense of community, as something that occurs only when students are physically co-present, new technologies have forced a reexamination of conventional definitions. It has been argued that the definition of communication may be better viewed as elocutionary co-presence, and the definition of community may be better viewed as people with common interests, and not simply physical common ground (see also Jazwinski, chapter 9, and Morahan-Martin, chapter 10). Requiring physical co-presence or proximity for meaningful interaction has become obsolete.

C. Conversational Pragmatics

Other concerns about the quality of computer-mediated versus classroom interactions stem from the fact that most on-line communication uses written language, while in "real" classroom discourse the predominant mode of communication is oral. On-line written communication suffers from a lack of social presence, and many metalinguistic and pragmatic cues, like gesture, expression, tone of voice, speaker volume, and distance, are unavailable. Further, the relationship between speaker and listener is altered, as consent or commitment on the part of the listener cannot be assumed, as it may be in face-to-face interaction. But some

advantages of on-line, written dialogue also can be noted. The discussion is often preserved in its entirety, whereas oral conversations are not usually preserved. The transmission of normal conversation is fleeting and therefore not entirely available for reflection either by the participants or by others who may be observing the discussion. While writing comments takes more time than speaking, this too may have its advantages as writing may require more effort, more planning, and more reflection than speaking (Joinson & Buchanan, chapter 11). Further, although on-line written discussions are frequently asynchronous, taking away the immediacy of the experience of participating in a group activity or community, many students may find the flexibility and lack of time pressure more convenient and conducive to deep thought.

On-line communication resources not only allow the learner to be more flexible in terms of where and when they participate in a dialogue, they may also allow a wider audience to read and participate in a discussion. This potentially means that more perspectives can be entertained and a more heterogeneous sample may be encountered (see chapters 9, 10, and 11).

D. Depersonalization

A related concern to the loss of social interaction in virtual space is the loss of identity. By talking through a machine, many have feared that computers are a dehumanizing and depersonalizing influence, resulting in socially undesirable behaviors like "flaming." However, depersonalization or anonymity can also have its advantages. Anonymity, along with the flexibility of communicating when and where you have time to think about the issues in an asynchronous on-line community, also allows more voices to be heard — particularly, students who may feel marginalized in a live classroom setting, women, minorities, and others who fear voicing an unpopular opinion or perspective may be more likely to communicate in an on-line arena. Students may also be more inclined to respond truthfully about their behaviors or their confusions in an anonymous interaction. And, anonymity can allow students to try out new roles or ways of relating to others (Jazwinski, chapter 9).

III. DIFFERENCES BETWEEN WEB LEARNING AND TEXTBOOK LEARNING

Along with differences between face-to-face and on-line conversation, there are also a number of important contrasts between textbook learning and web-based learning, mostly centered around the greater interactivity of electronic text or hypertext.

A. Diversity of Resources

The most striking difference between print and electronic media is the number and diversity of resources immediately available to students, and the student's ability to choose what to read. Students using the web have easy access to a seemingly infinite number of resources, many of them including multimedia presentations. Sources can be obtained representing many heterogeneous perspectives on any one issue or topic. Further, less structure is imposed on the order in which information can be read, so students have choices about how to navigate through information.

B. Permanence

Britt and Gabrys (chapter 4) discuss a number of differences in the "permanence" of written and electronic texts. Electronic sources can be altered and updated easily. This leads to some problems, as the content of pages can change, making the location of referred to information unreliable, and links frequently die. Even the layout of content on pages is not static. This is particularly so on sites with scrolling text. Students may not be able to have a reliable memory for "where" they read something, and even if they do, the site may not be there tomorrow.

C. Authenticity

In chapter 4, Britt and Gabrys also note that printed and electronic text differ in the sense of their writer's presence or identity. Anonymity is a two-edged sword in electronic text, much as it is in electronic discourse. Anyone can publish an opinion on the web. This makes authenticating information more difficult than from physically published material. On the other hand, since adding to the web is so easy, students have the ability to add, edit, and alter sources that they find, thereby having creative opportunities to contribute to the web in a way not possible with hard copies of texts (Hammond & Trapp, chapter 8).

D. Motivation

At least for the present, the novelty of web-based materials may be quite motivating for some students. In addition, considerable evidence suggests that web usage can be addictive (see Morahan-Martin, chapter 10). Although problematic in some contexts, the addictive nature of the Internet may also be quite useful for education, as textbook reading is almost never characterized as addicting. One intriguing, though rarely mentioned, possible source of the motivating quality of the Internet is its uneven quality. When surfing the web, good "hits" are typically interspersed with more mediocre ones.

Thus, individuals may experience random intermittent reinforcement, which is well known to encourage prolonged engagement in the reinforced activity. One curious implication of this view is that at least in some educational venues, some variety in the quality of materials might actually enhance participants' motivation to explore more thoroughly.

E. Familiarity

Although the unique novelty of web-based materials is likely to enhance students' motivational qualities (at least for some of them) the lack of experience and familiarity associated with such materials may also have significant costs. Due in part to our experience with published media, paper has a number of advantages, especially in the way we are used to reading text. Given a choice, most people will print out an electronic document in order to really read it. On-line reading is usually more like skimming or browsing than studying. There is a sense of immersion that people get from reading a book that people rarely get in reading an on-line article. Thus, Wolfe suggests in chapter 5 that brief presentations of text are best in an electronic or hypertext environment, and this seems suited to a browsing style of reading.

F. Multimodality

In the design of the Dragonfly Web Pages (Wolfe, chapter 5), the texts are actually both short and illustrated. The capability for including images seems one of the most obvious advantages of on-line text. Images can make text more interesting, or more meaningful. Multimedia capabilities can allow for more authentic learning environments and simulations. They can convey abstract principles to students. Similarly, tools that allow students to create or manipulate their own images can allow readers to solidify their understanding and make their abstract notions more concrete and available to others. At the same time, however, such images can result in overload, and can distract readers from important points in text (Reyna et al., chapter 2). Previous research advises that images and visualization tools like animation must be included judiciously when relevant for conceptual learning, or else the reader suffers too many demands on their attention (Harp & Mayer, 1998; Hegarty, Quilici, Narayanan, Holmquist, & Moreno, 1999; Wright, Milroy, & Lickorish, 1999).

G. Flexibility versus Linearity

Perhaps the feature of the web that has prompted the most techno-optimism, in terms of its educational promise, is its capacity for flexible, dynamic, and adaptive presentation of information. There are a number of

theorists who have suggested that the web represents a tremendous educational resource for students, especially because of its potential for self-directed and active learning. A reader may choose his or her own path through web documents, reaching beyond the confines of the narrow narrative path that traditional text must follow. With such flexibility the reader experiences the joys of discovery-based learning. Self-directed control of what is read next leads to more intrinsic interest in content, more motivation to learn, and more excitement in the learning process. Further, the reading of texts in an unstructured way requires students to generate their own sense of connection and coherence between the documents. Encountering isolated texts, instead of a single text, can in itself prompt more active processing of information, more reflection, and the recognition of novel connections between ideas. Thus, there are a number of romantic notions that hypertext may engender deeper thought and richer learning as learners navigate on paths of their own choosing through a network of ideas, and must actively construct meaning for themselves.

However, as many studies of text comprehension have demonstrated, the linear structure of a text serves an important function, as it gives the reader a way to navigate through information and a starting point for developing a representation of the text. While students may enjoy the empowerment and freedom of having their own navigational control, navigational choices may also overload the capabilities of the reader in some cases. While the linear structure inherent in most lectures and texts limits readers, it also supports their understanding. Access to an infinite number of choices gives readers the agency to follow their own interests and may allow for more intrinsic motivation to enter the learning context, but unless the reader has a great deal of knowledge about the content, or a good understanding of the web space, a lack of organizational structure can often leave students "lost in hyperspace." A major question for web-based education is how to support both navigation and representation while allowing for flexibility.

In chapter 7, Riva suggests an intriguing and important new tool, shared hypermedia (SHY), that may provide a good balance of student-initiated navigation and guidance. Shared hypermedia applications may be used to allow learners to choose their own destination, preserving their sense of agency and allowing them to follow their own interests. But, once at a site, students may be guided through the hypertext links and web pages in an instructor-led tour. This may provide enough support for learners to get their bearings within a site and get a sense of the important features and concepts, so that they can get the most out of their web experience.

A second way in which the dynamic, flexible, and adaptive nature of on-line text can benefit learning is through the inclusion of interactive features that force students to refine their understanding of the subject matter. There has been a great deal of converging evidence demonstrating that activities that require learners to engage in active, constructive, and integrative tasks

lead to the best understanding of the subject matter (e.g., Bereiter & Scardamalia, 1996; Brown & Campione, 1996; Kintsch, 1994; Mayer, 1989). Web environments in particular can be designed and used to capitalize on these modes of learning. Giving students an environment where they actively manipulate information and construct their own conceptual representation of the subject matter seems to be the best way to capture the benefits of new technologies. Wolfe's Dragonfly Web Pages provide an excellent example of this kind of learning opportunity.

H. Customization

A final advantage of the dynamic, adaptive, and flexible nature of on-line presentation is the capacity for all kinds of text, from short warnings to whole sites, to be customized to the specific needs of readers. By using short on-line assessments, as suggested by Joinson and Buchanan (chapter 11), readers can be directed to sites that are tailored for their needs. Anderson (chapter 3) has suggested a number of personality and cognitive styles that may benefit from sites with particular features. Alternatively, all students may be provided with a single version of a site, but prompts may provide individual-specific feedback or scaffolding in response to certain behaviors. Students can receive feedback, prompts, or background information when an intelligent program recognizes certain patterns of use, or when wrong or right answers are recognized in on-line exercises. They can thus be tutored on necessary skills, like how to perform effective searches. Dynamic processes can also be used to model expert behavior, and as mentioned in the SHY example, students can be guided through sites by a teacher, or directed to important concepts by links or hints.

It has long been held that there is an optimal match between reader and text that may result in the best learning outcomes. The adaptive nature of web presentation holds the promise to fit readers to ideal conditions by screening readers and altering the information that is presented to them in a manner suited to their abilities. Or alternatively, web environments could support readers into using behaviors and styles that are more appropriate for on-line learning.

IV. FUTURE RESEARCH

The chapters here have done a good job of laying out many possibilities for virtual education. What remains, as many have noted, is the empirical investigation of exactly which features of face-to-face and virtual learning environments lead to the best educational outcomes in which circumstances. Few of the chapters here included a formal assessment of learning outcomes or other evidence that the ideas they have suggested will actually translate

into better student learning. Unfortunately, this is also the case in the broader literature, as few studies on educational use of electronic text or multimedia resources have actually shown significant learning gains (Chen & Rada, 1996; Clark, 1983; Dillon & Gabbard, 1998; Landauer, 1995). In fact, few have measured learning outcomes in controlled experiments. While multimedia resources hold great promise as an educational tool, it is clear that putting a student in front of a computer screen, a web browser, or any other multimedia format does not necessarily lead to better learning. There are a wide range of outcomes that can be evaluated.

The most frequently reported outcome measure in the educational computing literature is preference or liking for a learning experience. While there need be no direct relation between enjoyment and actual learning, to the extent that preference measures predict whether students will return to the web in the future, and therefore have the opportunity to learn new content or skills, then liking seems to be an important variable to consider. Similarly attitudes toward technology may be an important mediating or outcome variable. More to the point, however, it is important to test actual effects of web-based education on students' learning of skills or content. This can be achieved by observing the use of skills. For example, have students learned to seek out the source of information, or do they formulate good Boolean searches? Learning can also be assessed through the obser- vation of behaviors. For example, do students spontaneously engage in higher-order mental processes, and do they look for contradictions? Learning can also be tested through the assessment of students' memory for facts or their understanding of concepts from web-based materials.

The previous chapters have suggested how the World Wide Web can support many student behaviors (i.e., interacting, active processing, sourc- ing) that may be important for better learning. Observation of which behav- iors may be promoted by technology is an important first step. And all of the ideas mentioned here seem theoretically promising. We hope that the inter- active game-like activities found on the Dragonfly Web Pages will lead to more active processing of information and therefore better understanding of the scientific concepts. We would hope that students who learn to consider the source of documents in their use of the web will develop a better under- standing of the subject matter than students who accept information unquestioningly. We hope that readers who use "deeper learning strategies" will have better retention of what they read. We would also hope that students who receive feedback as they are learning will do better in a course. However, all of these practices need to be tested empirically, as sometimes the best practices are not intuitively obvious.

Under many conditions, learning from multimedia or electronic text can certainly lead to poorer learning than occurs from traditional formats (Mayer, 1997; Rouet, Levonen, Dillon, & Spiro, 1996; Wiley & Voss, 1999). There may be nothing in a given medium per se that ensures better learning.

However, there may be certain learning behaviors that one medium engenders better than another. The key for educational research is to figure out how to allow students to gain the most from each medium. Technology offers us a number of reasons to be excited about its potential, and there are a number of theoretical approaches that lend themselves to implementation in electronic environments, which in turn may lead to better learning. Interactive tutors can supply important pieces of domain knowledge and guidance through complex reasoning and problem-solving exercises. This may allow novice learners to complete tasks that they would otherwise be incapable of performing. Computer environments can simulate real-world situations and help students ground their learning in concrete problem-solving contexts. Further, multimedia environments can provide students with good models of difficult or abstract concepts — through illustrations, diagrams, or animations — that they would have trouble visualizing through text alone. An important goal of the educational computing literature needs to be to determine which specific instructional contexts allow for effective educational uses of web resources.

In a good example of the kind of study that is needed, Reyna et al. (chapter 2) describe how the effects of on-line presentation can be assessed. Among three types of graphic presentations that accompanied a unit on heart medication, Reyna et al. found that students who saw a balloon-like diagram were more likely to use the important concepts of "flow" and "pressure" (and use them correctly) in their discussion of their problem solving. This protocol analysis suggests that students in this condition were acquiring the key concepts of the unit more so than students in the two other conditions.

In a similar kind of analysis, Wiley and Voss (1996) investigated the advantages of learning from a multiple-source environment over a textbook-like environment. Using Sourcer's Apprentice, the environment designed and described by Britt and Gabrys in chapter 4, students were presented source texts about the Irish Potato Famine either as eight separate documents in a web site or as a textbook chapter that contained the same exact text in a single document. Thus, the presentation format, but not the content of the information, was varied. Further, the task that students were assigned to do as they read varied. Some students were asked to write an argument about what produced the changes in Ireland's population, while others were asked to write a narrative. The essays that the students wrote were analyzed, and when students wrote arguments from multiple sources their essays contained more causal terms than those in the other three conditions. Students in the multiple-source argument condition were also more likely to integrate information from different sources. These findings suggest that students were developing a better understanding when they had multiple sources and a task that required them to integrate the sources. Like the findings in the protocols from the Reyna et al. study, this analysis

suggested which specific design features and uses of computerized learning environments lead to the best learning outcomes.

Although both of these studies used discourse analysis to probe the kinds of learning that occur in different learning contexts, there are also simpler ways of assessing whether deeper understanding is occurring. Verification tasks, although usually used for more superficial measures of memory, like retention of facts, can also be designed to test deeper levels of learning. Royer, Carlo, Dufresne, and Mestre (1996) have described how learning measures that tap multiple levels of representation can be developed. In a follow-up study, Wiley and Voss (1999), using the procedure of Royer et al., created three kinds of verification tasks to assess the kinds of learning that were occurring from the Potato Famine web site. In one task, students were simply assessed on sentence verification. In a second task, students were tested for inference verification and were asked to indicate whether a point seemed to follow from what they read. In a third test, a principle identification task, students were asked to rate the similarity of potential conceptual analogies to the cause of the Potato Famine. In this task, students were given four examples, one which was similar neither on the surface nor at a conceptual level, one which was similar at both the surface and the conceptual levels, one similar only on the surface level, and one similar only on an underlying conceptual level. The tests of most interest were those thought to reflect better understanding of the content: (1) good performance on the inference task, and (2) the recognition of the deep analogy. Consistent with the essay analyses, students who read from multiple sources and wrote an argument had the best performance on the inference and analogy tasks, indicating that they had indeed gained a better understanding of the subject matter.

Interestingly, students in the textbook and narrative condition of the Wiley and Voss (1999) experiment actually had the best performance on the verbatim memory task (the sentence verification task). This result is similar to that of others (e.g., Reyna et al., chapter 2). This discrepancy between verbatim memory and comprehension measures emphasizes the importance of including both retention and transfer outcome measures as learning assessments when measuring learning from the web. While multimedia and other computerized learning environments can be used to support more active learning, and potentially better understanding of the subject matter, students may actually have poorer surface or verbatim memory for information when compared to textbook or lecture-based learning.

Other examples of measures that test for understanding come from the studies of Mayer and his colleagues, who have investigated how images should be incorporated in text for effective learning (e.g., Harp & Mayer, 1998; Mayer, 1989). In these studies, he tests for "transfer" by giving a few short problem-solving tasks after students see units on scientific topics such as how brakes work and why lightning occurs. These problems involve

troubleshooting questions, questions about how machines could be made more reliable and effective, and questions about how effects could be counteracted or changed by hypothetical factors. In a similar series of studies, Hegarty has developed multiple choice questions that relate to the creation of an accurate causal model about the mechanisms that are studied (Hegarty et al., 1999). Similarly Kintsch and his colleagues have employed sorting tasks, where an understanding of the subject matter leads students to sort ideas into piles based on underlying concepts instead of more superficial associations (i.e., Mannes & Kintsch, 1987). These examples of verification, problem solving, and sorting tasks are included to show that detailed protocol analysis, or time-consuming analysis of navigation logs, while an important source of information on how students use and learn from the web, are not the only means available for empirical evaluation. These shorter tasks may be easier to administer and will also yield important and much needed evidence about effective uses of the web in teaching and learning.

V. THE WEB AS AN EXTENSION OF THE MIND

Although much research is clearly needed before we will be able to flesh out the full pedagogical implications of the web, it seems safe to predict at least one dimension that is certain to play an important role is in shaping the way we think about thought itself. The web has more general implications for cognition, particularly as a metaphor for the mind. As Riva notes in this volume, various researchers have suggested that the web represents something akin to an "interbrain" in which "cognitive activities are increasingly being performed in networked contexts which, to varying degrees, are undeniably virtual" (Riva, chapter 7). Such parallels between the net and the mind raise intriguing possibilities about the potentially unique learning value that a medium with such a close affinity to the mind might have. Importantly, however, Hammond and Trapp (chapter 8) point out that we must be very cautious in assuming that the net, by virtue of the fact that it shares some similarities with brains, should be assumed to lead to enhanced learning. Hammond and Trapp refer to this false conclusion as the "homeopathic fallacy," that is, that "the analogous network-like structures of both the web and the central nervous system (or, perhaps more plausibly, between the web and the associative structure of memory) somehow enhances a more direct transfer of information from computer screen to the mind." As they aptly put it, "to draw the inference that this similarity is a contributor to effective learning is a little like claiming that porridge is good for learning because it looks like the gray matter of the brain" (chapter 8). Indeed, the fact that to date there have been few studies that unequivocally demonstrate that the web enhances learning over traditional methods illustrates

that the web's parallels to the mind do not make it an unequivocally superior learning medium.

Although the parallels between the web and the mind have been thoughtfully considered in discussions of the web's potential pedagogical value, the mind–web metaphor has been largely ignored in traditional cognitive psychology theories, which still typically rely on the digital computer metaphor (Roediger, 1980). As Ulric Neisser observed, "models of the mind always follow the latest advances in gadgetry" (1982, p. 7). In the third century B.C., Plato likened the mind to a wax tablet. In the 17th century, Descartes fashioned many mental processes after the complex clock robots that were popular at the time. In the early 20th century, the mind was likened to a telephone switchboard, and of course, in the later part of the 20th century the mind was thought to resemble a digital computer. Given the relentless mapping between recent technology and the mind, it simply stands to reason that cognitive models will come to explore the parallels between the mind and the net.

In fact, when we consider the viability of the web as a metaphor for the mind, it is certainly far more apt than a wax tablet, and potentially provides a useful alternative or supplement to the still reigning (albeit with less strength than it once had) digital computer metaphor. There are certainly some significant parallels between navigation through the cyberspace of the World Wide Web and navigation through the mental space of the mind. In both cases, sophisticated search engines sort through a huge network of information to provide associations that are most relevant to the target information. Admittedly web search engines are not as effective as human search engines (yet), but perhaps as the parallels between the two become fleshed out, development of web search engines may become increasingly informed by our understanding of cognitive search mechanisms. Once one has landed on a particular location in search space (be it mental space or cyberspace), other associations present themselves. In both cases these associations are not random but rather are conceptually related in important respects to the target location. Just as we can choose which subsequent web sites to click on, so too do we move through mental space. Like lists from a search engine, multiple ideas often come to mind and we can "decide" which idea to pursue.

The web metaphor also nicely illustrates the relationship between conscious and unconscious processes. When we search our minds, we can control the basic topics that we choose to consider, and when alternatives come to mind we can choose between them. However, we have relatively little insight into the specific search processes that bring particular thoughts to mind, nor do we have much control over the nature of the specific thoughts that arise. Similarly, with the web, we can choose our search terms, but we cannot control what specific hits will come up. We can then choose what locations to visit, but we cannot control what those pages will look like.

Another useful attribute of the web analogy, in contrast to many of the other metaphors of the mind that have been introduced over the years, is that it includes a first-person perspective. We cannot imagine what it "feels" like to be a computer, and so the computer metaphor for the mind has always been hard to relate to. In contrast, it is easy to think about surfing one's internal mental space in a manner that resembles the way in which we search cyberspace because in both cases there is "someone" participating in the search. There are a number of exciting potential implications of the parallels in the subjective experience of searching mental space and cyberspace. First, it suggests that expertise in understanding how to navigate through the web may provide individuals with a natural and potentially rich knowledge domain for conceptualizing human information processing more generally. In short, the web metaphor of the mind may provide a particularly intuitive model of the mind that readily relates to a domain of universally increasing expertise.

In addition to providing a rich and accessible metaphor for conceptualizing human thought, the parallels between the web and the mind may actually influence the manner in which individuals ultimately learn to search through mental space. Surfing the web introduces a variety of formal demands for optimum searching. One needs to know how to define searches appropriately, how to assess the promise of the alternatives that arise, and to how maintain a general representation of what has been learned. All of these skills may similarly apply to searching mental space, suggesting that some of the information searching skills acquired through surfing the net might ultimately come to be internalized and used in searching one's own mind. Although provocative, the notion that individuals may come to internalize the information search skills that they use in exploring the web is consistent with a general view, originally articulated by Vygotsky (1978) and subsequently fleshed out by others (see Riva, chapter 7), that internal thought is a reflection of the social situations that individuals encounter. If, as a consequence of human interactions, individuals come to develop an internalized mode of thought that initially resembles something akin to social conversations, then it is not all that far-fetched to suggest that experiences on the Internet might result in a greater specialization of the modes of operation (such as defining problem searches and conceptualizing hyperlinked representations) that are particularly emphasized on the web. Indeed such claims are testable as they suggest that a relationship may be found (both correlationally and through extensive training) between individuals' experience with and proficiency in searching the web and their ability to search thoughts in their own minds. It may well turn out that although the web is not uniquely effective in getting information into our minds, it is uniquely suitable for encouraging the deft navigation of ideas within the complex representational network that constitute our personal "mental webs."

References

Bereiter, C., & Scardamalia, M. (1996). Rethinking learning. In D. R. Olson & N. Torrance (Eds.), *The handbook of education and human development: New models of learning, teaching and schooling* (pp. 485–513). Oxford: Blackwell Publishers.

Brown, A. L., & Campione, J. C. (1996). Psychological theory and the design of innovative learning environments: On procedures, principles, and systems. In L. Schauble & R. Glaser (Eds.), *Innovations in learning: New environments for education* (pp. 289–325). Mahwah, NJ: Erlbaum.

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

Clark, R. E. (1983). Reconsidering research on learning from media. *Review of Educational Research*, 53, 445–459.

Dillon, A., & Gabbard, R. (1998). Hypermedia as educational technology: A review of the quantitative research literature on learner comprehension, control and style. *Review of Educational Research*, 68, 322–349.

Harp, S., & Mayer, R. (1998). How seductive details do their damage: A theory of cognitive interest in science learning. *Journal of Educational Psychology*, 90, 414–434.

Hegarty, M., Quilici, J., Narayanan, N. H., Holmquist, S., & Moreno, R. (1999). Multimedia instruction: Lessons from evaluation of a theory-based design. *Journal of Educational Multimedia and Hypermedia*, 8, 119–150.

Kintsch, W. (1994). Text comprehension, memory, and learning. *American Psychologist*, 49, 294–303.

Kraut, R., Patterson, M., Landmark, V., Kiesler, S., Mukophadhyay, T., & Scherlis, W. (1998). Internet paradox: A social technology that reduces social involvement and psychological well being? *American Psychologist*, 53, 1017–1031.

Landauer, T. (1995). *The trouble with computers*. Cambridge, MA: MIT Press.

Mannes, S., & Kintsch, W. (1987). Knowledge organization and text organization. *Cognition & Instruction*, 4, 91–115.

Mayer, R. (1989). Models for understanding. *Review of Educational Research*, 59, 43–64.

Mayer, R. (1997). Multimedia learning: Are we asking the right questions? *Educational Psychologist*, 32, 1–19.

Neisser, U. (1982, January). *On the trail of the tape-recorder fallacy*. Paper presented at the symposium, "Influence Of Hypnosis And Related States On Memory: Forensic Implications," American Association for the Advancement of Science, Washington, DC.

Pew Internet and American Life Project (2000, May). *Tracking on-line life: How women use the internet to cultivate relationships with family and friends* [on-line]. Available (May 11, 2000): http://www.pewinternet.org/

Roediger, H. L. (1980). Memory metaphors in cognitive psychology. *Memory & Cognition*, 8, 231–246.

Rouet, J. F., Levonen, J., Dillon, A., & Spiro, R. (Eds.) (1996). *Hypertext and cognition*. Mahwah, NJ: Erlbaum.

Royer, M., Carlo, M., Dufresne, R., & Mestre, J. (1996). The assessment of levels of domain expertise while reading. *Cognition & Instruction*, 14, 373–408.

Vygotsky, L. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.

Wiley, J., & Voss, J. F. (1996). The effects of "playing" historian on learning in history. *Applied Cognitive Psychology*, 10, 63–72.

Wiley, J., & Voss, J. F. (1999). Constructing arguments from multiple sources: Tasks that promote understanding and not just memory for text. *Journal of Educational Psychology*, 91, 1–11.

Wright, P., Milroy, R., & Lickorish, A. (1999). Static and animated graphics in learning from interactive texts. *European Journal of Psychology of Education*, 14, 203–224.