John M. Slatin

Reading Hypertext: Order and Coherence in a New Medium

The basic point I have to make is almost embarrassingly simple: Hypertext is very different from more traditional forms of text. The differences are a function of technology and are so various, at once so minute and so vast, as to make hypertext a new medium for thought and expression—the first verbal medium, after programming languages, to emerge from the computer revolution. (The computer has spawned new media in the visual arts and music as well.) As a new medium, hypertext is also very different from both word processing and desktop publishing, two other computer technologies which have had an enormous impact on the production of texts. Both word processing and desktop publishing have as their goal the production of conventional printed documents, whereas hypertext exists and can exist only online, only in the computer. A new medium involves both a new practice and a new rhetoric, a new body of theory. I hope this essay will serve as a step in that direction.

The first requirement for a rhetoric of hypertext is that it must take the computer actively into account as a medium for composition and thought—not just as a presentational device and not simply as an extension of the typewriter. This means different things, depending on the level of abstraction at which one stands. At some levels, for example, one has to deal with data structures, with the construction of knowledge (if the term even applies) within and to the hard- and software of the computer as well as the construction of knowledge by and for humans. It means, also, that the rhetoric itself must be abstract, like Wallace Stevens's supreme fiction, in order to permit movement up and down the ladder of abstraction and to permit the articulation of principles that will enable practice. By the same token, the rhetoric of hypertext will have to be capable of change, for it is tied to a still immature (perhaps perpetually immature) technology which is itself changing at an exponential rate.

On the assumption, then, that description is the first step toward theory (see Bateson), I will contrast hypertext with more traditional text. I will focus on the assumptions each makes about what readers do and the ways in which those assumptions about reading affect the author's understanding of composition. For the purposes of this discussion, taking the computer into account means that we

John M. Slatin, associate professor, is Director of the English Department's Computer Research Lab at the University of Texas at Austin. He is the author of *The Savage's Romance: The Poetry of Marianne Moore* (1986) and several earlier essays on hypertext.

have to find ways of talking about documents that have multiple points of entry, multiple exit points, and multiple pathways between points of entry and exit points. Moreover, we must find ways to talk about the still more exciting kinds of activity fostered by this proliferation of possibilities: I mean interactive reading and its more or less inevitable concomitant, interactive writing, or co-authorship.

Widespread literacy is a comparatively recent phenomenon—that is to say that in Western societies such as those of England, Europe, and North America, general literacy is at best a couple of hundred years old. But Western culture was a print culture long before the coming of general literacy, and the text environment we are all familiar with is the product of fully mature, highly stable manuscript and print technologies which have been in place for many centuries. Principles and strategies for effective written communication are therefore based on assumptions about readers and reading. It will be helpful to consider these assumptions briefly before going on to discuss the different assumptions embedded in the design of hypertext and hypermedia systems.

The assumption that reading is a sequential and continuous process is the foundation on which everything else rests (see, for instance, Stanley Fish). The reader is expected to begin at a clearly marked point whose appropriateness has been determined by the author—usually with considerable effort: one of the hardest moments in any writing project is to figure out where to start—and to proceed from that beginning to an ending which is just as clearly marked and which has also been determined by the author in accordance with his or her understanding of the subject matter and the reader. The reader's progress from the beginning to the end of the text follows a route which has been carefully laid out for the sole purpose of ensuring that the reader does indeed get from the beginning to the end in the way the writer wants him or her to get there.

All but the most naive and inexperienced writers recognize that all but the most naive and inexperienced readers inevitably and rightly make inferences about what's going to happen next, on the basis of what they have already read—not only in the current text, but in other texts resembling it. The reader's perception of the predictability of a given text is an important factor in his or her qualitative evaluation of the text.

Prediction operates on a number of different levels and is determined by different things at different levels of abstraction. The predictability of a given text is a function of the relationships among phenomena at microscopic, macroscopic, and metatextual levels of abstraction. At the microscopic levels, the reader's ability to predict the course of the text from moment to moment is a function of such factors as paragraphing, sentence length, complexity of phrasing, vocabulary, and so on—that is, the factors that are evaluated in producing a so-called "readability index." Indeed, one might descend even further down the ladder of abstraction and argue that prediction takes place at the graphemic or phonemic levels as well (see Fry). At the macroscopic levels, the reader is aware of such things as general subject matter, topics and subtopics, and the structural devices organizing the text as a whole—sections, chapters, and subchapters, and so forth. At this level the reader is also at least subliminally aware of such things

as how much more material there is to read. On what we might call the metatextual level, the reader makes inferences about the text as a whole, based on his or her understanding of the larger context to which he or she regards the text as belonging. These inferences are often implicit, or else they may take the form of mental or marginal annotations. In any case, they are outside the text and separate from it; they are integral to hypertext, as I shall explain later on.

The end product of the writing process is a text written or printed on paper, and then, often, sewn or glued or otherwise bound between covers. This is obvious, I know, and yet it needs to be said. Every writer has had the awful experience of opening a book or article hot off the press, only to stare in horror at a glaring error in a crucial passage that somehow escaped the most agonizing scrutiny. The fixity of the printed text as an object in physical space makes the text as an object in mental space seem equally stable and fixed. Or at least that's how we tend to want it. As Richard Lanham has said, "It was establishing the original text that the Renaissance scholars thought their main task, and generations of textual editors since have renewed their labors. The aim of all this was to fix the text forever" ("Convergent Pressures" 4). The continuing controversy over Hans-Walter Gabler's edition of Joyce's *Ulysses* (1984; see Rossman) makes abundantly clear just how intense the desire "to fix the text forever" can be. Gabler's re-conception of the editing process may have occurred independently of his decision to use the computer (Groden 29), but the controversy over his "synoptic" version of *Ulysses* offers a clear illustration of the way computers will revolutionize our understanding of text. Every word in Gabler's synoptic text was written by Joyce himself—and yet the final "reading text" is a text no one ever wrote—it had never existed prior to its publication. (In this sense, Gabler's *Ulysses* resembles Benoit Mandelbrot's fractals, in which recursive mathematical formulae are graphically plotted to produce visual structures that, while in some cases resembling phenomena in the day-to-day world, have no counterpart in that world.) Gabler's is a simulated *Ulysses*, like the *Don Quixote* that would be produced by Borges' Pierre Menard if Menard were real and if he had a computer. This is what becomes of the work of art in an age of electronic reproduction. Text is always mutable, always subject to inadvertent error and deliberate change, and it has to be coerced into standing still. That's why publishers charge you money if you make too many changes in a text after it's been typeset.

For all these reasons—because a text looks like a permanent thing, because readers expect to begin at the beginning and end at the end and to know which is which (that's why students so often begin the last paragraph with "in conclusion"), because readers expect to get from beginning to end via a clearly-marked route—sequence is of paramount concern to a writer. Much of his or her effort goes into figuring out the correct sequence for the material that's going to be presented. The writer's job in this context is to contrive a sequence that will not only determine the reader's experience and understanding of the material but will also seem to the reader to have been the only possible sequence for that material; you want it to seem to have been somehow inevitable.

Of course this inevitability has a good deal to do with the issue of predictability I raised earlier. Readers have to be able to predict what will come next, at least up to a point, or they start to feel lost, which makes them start to feel nervous, which makes them want to put down what they're reading and go watch a football game or something—at which point the writer has failed miserably. But the flip side is just as bad, and the end result is going to be the same.

Writing that's too predictable is governed by a presupposition succinctly expressed by a certain hotel chain's ad campaigns. The presupposition that there should be "no surprises" may be fine for hotels. But it becomes a fundamental conceptual error where writing is concerned. The informational value of a given document is not simply a function of the quantity of data it presents or the facts it contains. At one level of abstraction, what we call information may indeed consist in numbers, dates, and other data, other facts. But as Gregory Bateson says, "All receipt of information is necessarily receipt of the news of *difference*" (32). At a somewhat higher level of abstraction, therefore, none of these data can be considered information until they have been contextualized, arranged in such a way that both the significant differences and the significant relationships among them may become apparent to the intended reader. In Christopher Dede's terms, this is when information becomes knowledge. In other words, as literary artists and their readers have always known, there can be no information without surprise.

Rhetoric typically has little to say about the physical processes by which a text is brought into being. Or I could put it even more strongly and say that rhetoric has traditionally been indifferent to the technology of communication. One reason for this indifference is that the technology is so mature that it's simply taken for granted, that it is essentially invisible *as* technology. There was a point in history, of course, when writing itself was a radically innovative technology and was regarded as such, as Eric A. Havelock, Father Walter J. Ong, and Richard Lanham have shown us. The computerization of writing has similarly made the technology itself highly visible, especially in the cases of desktop publishing and hypertext/hypermedia. By contrast with traditional text, hypertext and hypermedia depend upon an emergent technology which is still immature and still subject to radical transformation; indeed, all indications are that accelerating change is an inherent characteristic of this technology. It may *never* stabilize. Thus rhetoric for hypertext cannot afford to disregard the technological substrate upon which composition and reading depend.

There are many continuities between conventional text and hypertext. Anyone involved in creating a hyperdocument will still have to worry about the problems I've outlined so far. But hypertext is a very different kind of beast than a conventional text, and creating a hyperdocument poses some very different problems as well. The remainder of this paper will concentrate on those differences and their implications.

First of all, the hyperdocument may well contain material from different media such as text, graphics, video, and sound. While this is an important factor, I don't think it's decisive. After all, printed books often contain text, line

drawings, tables of data, reproductions of visual images, and so forth—though of course they cannot manage full motion video or sound. Besides containing different types of materials than those to be found in printed text, the hyperdocument is likely to contain considerably more material than a printed book. Again, this is not a decisive difference in itself: encyclopedias also contain an enormous quantity of material. The quantity of material in a hyperdocument does pose problems, and it does make for complexity. But the greatest difference between text and hypertext is not in the relative quantity of material each form handles: it's in the technology that handles the material.

What makes all the difference in the world is the fact that hypertext exists and can exist only in an online environment. This is crucial, not just because it substitutes monitors, keyboards, and mice for the customary physical apparatus associated with text—paper, books, pencils, and so forth. The fact that hypertext exists only in the online environment is crucial because, as Douglas Hofstadter says, "It is the organization of memory that defines what concepts are" (528). Hypertext uses machine memory in a way that has no analogue in the traditional text environment, where composition relies on the organization of human memory. It is the organization of memory in the computer and in the mind that defines hypertext and makes it fundamentally different from conventional text.

In such an environment, the problem is not simply to develop effective strategies for implementing well known and long established principles of effective communication. On the contrary, one of the chief functions of rhetoric in the hypertext environment is to discover the principles of effective communication and then develop ways of implementing those principles through the available technology.

The rapidly evolving technological environment makes hypertext possible by permitting the embodiment of a very different set of assumptions about readers and reading—and about thinking. These assumptions in turn form the basis for decisions made in the process of creating a hyperdocument.

Reading, in hypertext, is understood as a discontinuous or non-linear process which, like thinking, is associative in nature, as opposed to the sequential process envisioned by conventional text. Associative thinking is more difficult to follow than linear thinking. Linear thinking specifies the steps it has taken; associative thinking is discontinuous—a series of jumps like the movement of electrons or the movements of the mind in creating metaphor. This discontinuity is not fortuitous; rather, as Stewart Brand points out, it is a basic aspect of the digital encoding of information. Brand offers the illuminating contrast between the surface of a traditional phonograph album, with its continuous grooves, and the surface of a compact disc, with its distinct, discontinuous pits (18).

Reading in this sense has little to do with traditional notions of beginning at the beginning and going through to the end. Instead, the reader begins at a point of his or her own choosing—a point chosen from a potentially very large number of possible starting points. The reader proceeds from there by following a series of links connecting documents to one another, exiting not at a point defined by the author as "The End" but rather when he or she has had enough. Accordingly, the most common metaphors in discussions of hypertext equate reading with

the navigation or traversal of large, open (and usually poorly-charted) spaces. As Jeff Conklin has pointed out, because the hyperdocument contains so much material, and because relations between the components of the hyperdocument are not always spelled out, there is a significant danger that the reader will get lost or become badly disoriented.

The difficulty is compounded because hypertext systems tend to envision three different types of readers: the reader as browser, as user, or as co-author. The relationship between these three classes can be fuzzy and therefore difficult to manage. One function a rhetoric for hypertext will have to serve will be to provide ways of negotiating it.

The browser is someone who wanders rather aimlessly (but not carelessly) through an area, picking things up and putting them down as curiosity or momentary interest dictates. In this respect the browser is someone who reads for pleasure, with this important difference: there is no expectation that the browser will go through all of the available material; often the expectation is just the reverse. It is difficult to predict the browser's pathway through the material—and in fact it is less important to predict the pathway the browser will take than it is to provide a backtracking mechanism, what Mark Bernstein calls a Hansel-&-Gretel trail of breadcrumbs to allow the browser to re-trace his or her steps at will. (Of course this same mechanism is essential for readers in the two remaining categories as well.)

By contrast with the browser, the user is a reader with a clear—and often clearly limited—purpose. He or she enters the hyperdocument in search, usually, of specific information and leaves it again after locating that information. The user's path is relatively predictable, provided those who have created the hyperdocument have a sufficient understanding of the task domain. In these respects, then, the user resembles a typical student doing the assigned reading for a course. But there is also an important difference between the user and the student, which is most clearly recognized from the vantage-point of the author rather than the reader. The author(s) of a hypertext documentation system (e.g., a software product like Microsoft QuickBASIC 4.5™) will have met their goal when the user finds the information he or she needs and returns to the work in progress. However, the instructor designing a set of hypertext course materials may well not be satisfied with such an outcome. The instructor aims at a dynamic process, in which the student moves among three different states: from a user the student becomes a browser (and may then become a user once again); ultimately, he or she becomes fully involved as co-author. Thus what looks like a hierarchy of readers collapses.

One of the most important differences between conventional text and hypertext is that most hypertext systems, though not all, allow readers to interact with the system to such an extent that some readers may become actively involved in the creation of an evolving hyperdocument. Co-authorship may take a number of different forms—from relatively simple, brief annotations of or comments on existing material, to the creation of new links connecting material not previously linked, to the modification of existing material or the creation of new materials, or both. Both literary theorists (e.g., Wolfgang Iser, Paul Ricoeur, Stanley Fish)

and cognitive scientists like Jerome Bruner have talked for years about the reader's involvement in the construction of textual meaning. But hypertext's capacity for literally interactive reading and co-authorship represents a radical departure from traditional relationships between readers and texts. The implications of this departure from traditional relationships between readers and texts are enormous, both for the creative arts (see, for instance, Moulthrop; Ziegfeld) and for education: as many theorists now agree, understanding comes about when the mind acts upon the material. Marshall McLuhan's distinction between hot and cool media is relevant here—a cool medium being one that invites active participation, a hot one being one before which one sits passively. McLuhan was thinking, of course, about the difference between print and television, but one might argue that hypertext combines the heat and visual excitement of film, video, and television with text's cool invitation to participate (also Lanham, "The Electronic Word").

"Writing," in the hypertext environment, becomes the more comprehensive activity called "authoring." Authoring may involve not only the composition of text but also screen layout and other things that fall under the general rubric of interface design; it may also involve a certain amount of programming (as in Apple's HyperCard, where complex navigational and other processes are scripted by the stack's author). Perhaps most importantly, authoring involves the creation and management of links between nodes.

Ted Nelson, who coined the term "hypertext," defines hypertext as "nonsequential writing" (117). This means writing in which the logical connections between elements are primarily associative rather than syllogistic, as in conventional text. One implication of this is that the hyperdocument "grows" by a process of accretion, whereas the conventional document tends to have been winnowed out of a larger mass of material. That is, in preparing to write a conventional document, you almost inevitably assemble more material than you can possibly use; the closer you come to final copy, the more you find yourself excluding material that "doesn't fit" the subject as you've finally defined it. Hypertext, by contrast, is an inclusive medium. Thanks to the capability of creating nodes and links, material not linearly related to the point being discussed at the moment but still associated with that point may be placed in a node of its own and linked to other nodes as appropriate; the material need not be thrown away (see also Howard). In much the same way, no individual point of view need be excluded.

This inclusiveness makes it unlikely that any one individual will see all the elements making up the system. It also means that the hyperdocument is in fact a collection of possible documents, any one of which may be actualized by readers pursuing or creating links between elements of the system (Slatin, "Hypertext and the Teaching of Writing" 113).

The end product of the authoring process, the hyperdocument, is not a closed system, like a book; it is rather an open and dynamic system. The hyperdocument is an online system or network whose constituents are of two basic types, nodes and links. Nodes may consist of documents, images, or other materials electronically connected—linked—to one or more other documents or images.

Very likely, the different nodes will represent the work of quite a few individuals, who may have been working at different times and in different locations. Indeed, one impetus for the development of hypertext systems has been in the need to address exactly this issue among members of development teams (see Conklin; see also Trigg, Suchman, and Halasz; Suchman and Trigg). The development of protocols and procedures for co-authorship thus becomes an important issue. So does the development of procedures for moving through the system.

The reader's progress through a conventional text is governed by the arrangement of the material; the burden of prediction falls more heavily upon the reader than on the writer. This situation becomes considerably more complicated in hypertext. Given a system of discrete and interconnected nodes, the reader/user must decide which links to follow; in order to make that decision intelligently, he or she must be able to make reliable predictions about the consequences of particular choices. But the freedom of movement and action available to the reader—a freedom including the possibility of co-authorship—means that the hypertext author has to make predictions as well: for the author, the difficulty at any given moment is to provide freedom of movement and interaction, while at the same time remaining able to predict where the reader/user will go next. The most effective solution here, I think, will be to treat each node as if it were certain to be the reader's next destination. This is time-consuming in the short run, but in the long run probably saves time by creating a more readily usable system.

This brings us to the issue of linkage, the mechanism that creates the hyperdocument and allows the reader to move through it. Douglas Hofstadter has suggested that the perception of relatedness is a defining characteristic—perhaps *the* defining characteristic—of intelligent behavior. Hypertext embodies this idea, for everything in hypertext depends upon linkage, upon connectivity between and among the various elements in the system. Linkage, in hypertext, plays a role corresponding to that of sequence in conventional text. A hypertext link is the electronic representation of a perceived relationship between two pieces of material, which become nodes once the relationship has been instantiated electronically. That is, the link *simulates* the connections in the mind of the author or reader; and it is precisely because the electronic link is only a simulation that problems will arise.

The interdependency of links and nodes is such that it is impossible to talk about one without talking about the other. Thus the question of how to define a node leads to additional questions about linkage. These lead in turn to questions about structure (Slatin, "Hypertext and the Teaching of Writing") and coherence, and so back again to the issue of prediction.

A node is any object which is linked to another object. It may be as large as an entire book or as small as a single character; theoretically, it could be as small as a single pixel (picture element) on a display, though such extremes seem hardly practical. It may consist of a document or a block of text within a larger document; it may be a drawing, a digitized photograph, a (digitized) detail of a painting, a sound recording, a motion picture, or a scene—even a frame—from a motion picture.

There is no set answer to the question, how big should a node be? just as there is no set answer to the question, how long is a paragraph? Like the paragraph, the hypertext node is a way of structuring attention, and its boundaries, like those of the paragraph, are somewhat arbitrary. A node may contain a single paragraph; it may contain many; it may contain something else entirely. Jeff Conklin offers criteria for guidance in determining node size. To what extent, he asks, is the information in question so tightly bound together that (a) you always want to view it together; (b) you never want to take it apart; and (c) you rarely even want to reference parts of it outside of the context of the rest (38–40). In other words, a node is an integrated and self-sufficient unit; its size will be a function of the complexity of the integration. This in turn is contingent upon the author's perception of the nature of the material it contains and the relation of that material to other things in the hyperdocument.

The individual node, then, behaves in certain respects like a more conventional text. But the node is not just a self-contained unit. A node cannot, by definition, be entirely free of links—a node is a knot, is always embedded in a system—and that connectedness in turn gives the node its definition: "A node is something through which other things pass, and which is created by their passage" (Slatin, "Hypertext and the Teaching of Writing" 126).

A text becomes a node when it is electronically placed in relation to other materials (documentary or otherwise), which may (or may not) already contain links to other elements within the system. The difficulty here, of course, is that what are self-evident associations to me may not be even faintly apparent to you, and vice versa. This imposes an obligation on the author(s) of a hyperdocument which has no exact parallel for the writer of conventional text: the nodes must seem complete in themselves, yet at the same time their relations to other nodes must be intelligible. The problem of relationality here, as I argue elsewhere, is analogous to the problems of intertextuality confronting readers of, say twentieth-century poetry. This problem becomes increasingly challenging as the hyperdocument expands. Links exist for many different reasons—that is, to represent many different kinds of relationships between objects. The more links there are between the current node and other elements of the hyperdocument, then, the greater the necessity of identifying the attached material clearly—especially when the reader is allowed considerable freedom in choosing among the available links. (See Landow, "Relationally Encoded Links" and "Hypertext in Literary Education" for several pertinent rules of thumb.)

These identifiers carry an enormous burden. Indeed they are often asked to do the kind of explanatory work that ordinarily takes several sentences or paragraphs. Not surprisingly, there are several different methods of identifying link or node types. Some hypertext systems, such as MCC's gIBIS, use "typed nodes" (Conklin), while others—Xerox PARC's NoteCards, for example—employ "link types" (Trigg, Suchman and Halasz). These systems allow the coauthor creating a new node to choose from a list of pre-defined relationship categories, whose names then become part of the node or link.

HyperTIES, developed by Ben Shneiderman at the University of Maryland and distributed by Cognetics Corporation, offers a variation on this approach by

encouraging the author to compose a brief (two-line) description of the linked node. In HyperTIES, where reading is defined primarily as browsing and where browsing is completely separate from authoring, activating a link is a two-step process. The first step, clicking on a highlighted word or phrase, brings up the description of the attached node; the second step either brings up the attached node or returns the browser to the current screen.

The developers of Intermedia at Brown University have chosen a third alternative. Links belong to "webs" rather than to documents; webs are displayed onscreen as visual maps. An Intermedia document is first displayed as if it were freestanding; then, when the user opens a web, the links belonging to that web are displayed (Garrett, Smith, and Meyrowitz 1986). This approach allows an individual node to be placed within multiple frames of reference, or "webs," while avoiding both the screen clutter and the mental clutter that can accrue so easily when multiple links radiate to and from a node. Which links the user sees will depend upon which web he or she has elected to open.

As the Intermedia approach suggests, one way to address the question of how many links a node should have is by turning it into the question of how many links should be displayed at any moment. Research on memory suggests that we can hang on to between five and seven "chunks" of information at a time and that creating links between these chunks is a way to increase the effective size of the chunk. However, the number of chunks that can be retained decreases in inverse proportion to their size (see Kozma).

Probably no single method of identifying nodes and their relationships to other nodes is adequate to all needs; some combination will be needed. Choosing from a list of predefined relationships has certain advantages, since there is a strong likelihood that a newly created node will fall within an existing classification. However, the list must offer real choices without becoming so big as to make choice impossible. And co-authors may also wish to give a fuller description of the attached node than the list approach permits; something like the HyperTIES strategy becomes appropriate here. And when a given node has multiple links to and from other nodes, it may be advisable to use an Intermedia-style mapping strategy.

The approach you choose to the problem of identifying links and nodes will depend on several factors: your understanding of the ways in which the material is related; your sense of who your readers are (are they primarily browsers? users? co-authors?). Your sense of what you want those readers to do is especially important.

You don't have to worry about interactive readers when you write a conventional text—the only thing you want the reader to do is to go on to the next sentence. But in hypertext, where there are a number of possible "next sentences" or nodes for the reader to go on to, you do have to make some decisions about what ought to happen next. That is, do you care whether the reader (a) opens a specific node or sequence of nodes; (b) chooses more or less randomly from the available links; (c) creates a new node, linking it not only to the current node but also to such other nodes as the reader—now a co-author—deems appropriate?

If you want the reader to open a specific node or sequence of nodes, you can either try to influence the reader's course of action, for example by highlighting a "preferred pathway" through the material, or you can simply pre-empt the reader's choice by automating the sequence or hiding links you don't want the reader to pursue. (Though if you take that route you give up many of the advantages of hypertext, it seems to me.) Or if you don't have a preference about the sequence the reader follows, you may opt not to give directions, leaving the choice of which links to activate—or whether to activate any link at all—entirely up to the reader. If you want to encourage response—that is, if you want the reader actually to get involved as a co-author—you should say so somewhere on the screen and make it as easy as possible for the reader to change roles. In HyperCard, for instance, you can script a button to open up a text field with a date/time stamp and the name of the new co-author; you can even let the co-author enter keywords to make later searches faster.

I've already said that the author of a hyperdocument has a hard time trying to predict where the reader will go from any given point. The reader who activates a link often has a hard time, too, because it can be so difficult to predict what the result will be. The more cryptic the link or node identifiers are, the harder it is for the reader to predict the results of activating a particular link. The harder it is to make such predictions, the greater the likelihood that the reader will simply opt out of the process in frustration. And even if the reader does go ahead there is no guarantee that he or she will know the place when he or she gets there.

The reader has to make several different kinds of predictions. First, he or she has to predict the kind of material he or she will encounter upon activating a link. It would be quite distressing, for example, to activate a link in the expectation of moving to a narrative explanation of some issue, if in fact the associated node contains only raw data in tabular form. Second, he or she makes some predictions about the content of the node: at the most general level, the reader makes some kind of assumption about the closeness of the relationship between the current node and the material linked to it.

The questions of link and node labeling which are obviously of such central importance here impinge on the issue of predictability at what I have been calling the macroscopic level. At the microscopic level, predictability revolves around such things as screen-design: typography, visual effects, layout of information on the screen, and so forth, all have an impact on the reader's ability to organize the material in his or her own mind, and thus on his or her ability to operate effectively within the hypertext environment. Typographical conventions are such that the reader of a conventional text has a pretty good idea what kind of object will appear next. The period at the end of a sentence leads one to expect a capital letter and the beginning of a new object belonging to the same class—a new sentence; the blank space at the end of a paragraph signals the beginning of a new paragraph; and so on. The signals in hypertext systems aren't nearly so clear.

Because the technology isn't mature enough yet to support a single set of conventions, each hypertext system has to develop its own conventions. For that reason, it is probably necessary to incorporate procedural discussions about these conventions into the hyperdocument itself; thus one element of the hyper-

document will be an ongoing critique of its own procedures. Participants might consider, for example, whether to assign specific fonts to individual co-authors (see Trigg, Suchman, and Halasz) so that their contributions can be readily identified; text styles might be assigned to specific node or link types; a question might always be italicized, for instance, while an explanation might be underlined; and so on. Color can be used for similar purposes in systems where color is available. So can visual effects such as wipes, dissolves, and zooms, which are available in HyperCard. These and other special effects can easily become distracting or even annoying, but if particular devices are consistently and intelligently used in association with particular node types, they can also function as more or less subliminal aids to prediction, helping the reader to perceive the hyperdocument as coherent.

We regard a conventional text as coherent to the extent that all the material it contains strikes us as being related in an appropriately direct way to the subject and to the author's thesis and arranged in the appropriate sequence. The perception of coherence in hypertext seems to me much more problematic, however, though I don't have time to do more than suggest what might be involved. Nor do I know enough to do more than that.

I think of hypertext coherence as appearing at the metatextual level—that is, at the level where the reader preceives what Bateson calls "the pattern which connects" (12). The pattern which connects is the organizing notion around which all the disparate elements of the hyperdocument revolve. (An author's feel for the pattern which connects plays a significant part in decisions about node size and linkage as well.) This can be a relatively straightforward thing—a given hyperdocument might contain all the materials generated during a particular design project, for instance. This metatextual level is perhaps best represented by a visual map of some kind, whose nodes would open up to map subordinate patterns. This map ought to be readily accessible from any point in the hyperdocument, which suggests that it might be "iconized" and placed at a consistent screen location. This sounds simple enough, perhaps; but it becomes problematic again when we remember that we're dealing with a fluid system and multiple participants, and we start to ask whose understanding such maps represent. Maybe there needs to be a facility to allow any user to create such a map, whether for private consideration only or for public use might be up to the reader/co-author.

Conceptually, hypertext has a place, I think, in any environment where it's necessary or desirable to bring together large, complex, highly diversified bodies of information in such a way as to emphasize their interconnectedness—especially if physical space is at a premium, as of course it is on board a space station or an orbital device, or in a control room—or, for that matter, in a classroom (see Slatin, "Text and Context"; Slatin, "Hypertext and the Teaching of Writing"; see also Jaffe and Lynch; Bourne et al).

Perhaps the greatest value of hypertext is its ability to link enormous quantities of material that, in a conventional text environment, would be kept separate, perhaps even in different buildings, so that things which someone perceives as being related do in fact become related. Hypertext is weakest when it comes to spelling out what these relationships entail. It is important to say this because

the techniques for explanation are quite highly developed within traditional rhetoric, and it would be a mistake to abandon them as outmoded.

Hypertext places different demands on both readers and authors than demands facing readers and authors of conventional text. The principal reason for this, in my view, is that hypertext is truly a new medium. Employing the full resources of technology to represent and correlate information, hypertext grants both readers and authors an unprecedented degree of freedom to arrange materials as they deem best, and it permits interaction between readers and authors to an unprecedented degree. In so transforming the methods of organization which have served traditional text for millennia, hypertext requires authors and system designers to find new methods of indicating relationships, representing and constructing knowledge, and achieving coherence.

Works Cited

Bateson, Gregory. *Mind and Nature: A Necessary Unity*. New York: Bantam, 1980.

Bernstein, Mark. "The Bookmark and the Compass: Orientation Tools for Hypertext Users." *ACM SIGOIS Bulletin* 9.4(1988):34–45.

Bourne, John R., et al. "Intelligent Hypertutoring in Engineering." *Academic Computing* 4.1(1989):18–20, 36–48.

Brand, Stewart. *The Media Lab: Inventing the Future at MIT*. New York: Viking, 1987.

Conklin, Jeff. *A Survey of Hypertext*. MCC Technical Report Number STP-356-86. Austin, TX: MCC, 1986.

Dede, Christopher. "The Role of Hypertext in Transforming Information into Knowledge." NECC. Dallas, 15–17 June 1988.

Fish, Stanley. *Self-Consuming Artifacts: The Experience of Seventeenth-Century Literature*. Berkeley: U of California P, 1972.

Fry, Dennis. *Homo Loquens: Man as a Talking Animal*. Cambridge: Cambridge UP, 1977.

Garrett, L. Nancy, Karen Smith, and Norman Meyrowitz. "Intermedia: Issues, Strategies, and Tactics in the Design of a Hypermedia Document System." Conference on Computer-Supported Cooperative Work. Austin, TX, 3–5 December 1986.

Groden, Michael. "Editing Joyce's *Ulysses:* An International Effort." *Scholarly Publishing in an Era of Change: Proceedings of the Second Annual Meeting, Society for Scholarly Publishing, Minneapolis, Minnesota June 2–4, 1980*. Ed. Ethel C. Langlois. Washington, D.C.: Society for Scholarly Publishing, 1981. 27–34.

Havelock, Eric A. *The Literate Revolution in Greece and Its Cultural Consequences*. Princeton: Princeton UP, 1982.

Hofstadter, Douglas R. *Metamagical Themas: Questing for the Essence of Mind and Pattern*. New York: Bantam, 1986.

Howard, Alan. "Hypermedia and the Future of Ethnography." *Cultural Anthropology* 3.3(1988):304–15.

Jaffe, Conrade C., and Patrick J. Lynch. "Hypermedia for Education in the Life Sciences." *Academic Computing* 4.1(1989):10–13, 52–57.

Kozma, Robert. "Designing Cognitive Tools for Computers." Hypertext Workshop. Intelligent Systems Branch, Air Force Human Resources Laboratory, Brooks AFB, TX, 23 February 1988.

Landow, George P. "Hypertext in Literary Education, Criticism, and Scholarship." *Computers and the Humanities* 23.3(1989):173–98.

———. "Relationally Encoded Links and the Rhetoric of Hypertext." *Hypertext '87 Papers*. Chapel Hill, 13–15 November 1987.

Lanham, Richard A. "Convergent Pressures: Social, Technological, Theoretical." Conference on the Future of Doctoral Studies in English. Wayzata, MN, April 1987.

———. "The Electronic Word: Literary Study and the Digital Revolution." *New Literary History* 20.2(1989):265–90.

McLuhan, Marshall. *Understanding Media: The Extensions of Man*. New York: New American Library, 1964.

Moulthrop, Stuart. "Containing Multitudes: The Problem of Closure in Interactive Fiction." *ACH Newsletter* 10(1988):1,7.

Nelson, Theodor Holm. *Literary Machines*. San Antonio, TX: Theodor Holm Nelson, 1987.

Ong, Walter J., S. J. *Orality and Literacy: The Technologizing of the Word*. London: Methuen, 1982.

Rossman, Charles. "The New *Ulysses*: The Hidden Controversy." *New York Review of Books* 25.19(1988):53–58.

Slatin, John M. "Hypertext and the Teaching of Writing." *Text, Context, and Hypertext*. Ed. Edward Barrett. Cambridge, MA: MIT P, 1988. 111–29.

———. "Text and Context: Reflections on the role of the Computer in Teaching Modern American Poetry." *Humanities and the Computer: New Directions*. Ed. David S. Miall. Oxford: Oxford UP, 1990. 123–35.

Suchman, Lucy A., and Randall H. Trigg. "A Framework for Studying Research Collaboration." Conference on Computer-Supported Cooperative Work. Austin, TX, 3–5 December 1986.

Trigg, Randall H., Lucy A. Suchman, and Frank G. Halasz. "Supporting Collaboration in NoteCards." Conference on Computer-Supported Cooperative Work. Austin, TX, 3–5 December 1986.

Ziegfeld, Richard. "Interactive Fiction: A New Literary Genre?" *New Literary History* 20.2(1989):340–72.