

42. [Introduction] Siren Shapes Exploratory and Constructive Hypertexts

The hypertext of the Web is not the hypertext imagined by Vannevar Bush (¶02), Doug Engelbart (¶08, ¶17), or Ted Nelson (¶11, ¶21, ¶30)—as reading these authors makes clear, the Web edition is much more limited. Understanding the limitations of the Web’s hypertext is not simply an occasion for complaint, however. It helps reveal the potential that still lies within the hypertext concept, untapped by mainstream new media. In the following essay, Michael Joyce gave a name to an important distinction between two types of hypertext environments—those that are “exploratory” and those that are “constructive.” His distinction maps onto significant differences between the environment in which we currently experience the Web and the ideas of early hypertext creators, while also usefully describing other areas of new media, helping reveal both limitations and opportunities.

According to Joyce, *constructive hypertexts* are those in the process of creation by the user/author. They are flexible representations of thoughts, stories, arguments, and everything else for which we use media. Exploratory hypertexts are former constructive hypertexts, now being experienced by a user/reader who is not an author of the work. In an environment such as those described by Bush, Engelbart, and Nelson, a user can freely move back and forth between the roles of author and reader, between the experiences of construction and exploration. A constructive hypertext environment provides something more than an electric textbook or workbook: it makes all exploratory hypertexts into material that the user/author can place within the space of the constructive hypertext they are always already creating. This process of creation, Joyce argues, holds much more promise for transforming education than does simple electric reading with links.

Constructive hypertext has largely remained an exciting activity on the periphery of education. But the tool that Joyce, J. David Bolter (¶47), and John B. Smith created to embody their goals of constructive, “topographic” hypertext creation has been very influential in another area: electronic literature. *Storyspace*, marketed now by Eastgate Systems, was, before the rise of the Web, the primary creative medium for hypertext literature, rivaled only by Apple’s Hypercard.

The Web has yet to include many of the features that Joyce associates with exploratory hypertext, such as the ability for users to create annotations and links—perhaps because no one has developed a revenue model to support offering such services over and above the current Web environment’s basic features. IBM offered an early service to allow users to add and share links between Web sites. This service, Aqvi, has been discontinued. Third Voice later offered a mechanism for making public annotations to Web pages—until the company folded. Link servers are now available from Active Navigation and other companies, but as an internal service for organizations, not a public facility. No major Web software even allows the user to create private links and annotations during reading.

Such facilities might be added to a browser such as Mozilla (perhaps along with alterations to a server such as Apache) by enterprising individuals. The more essential question is whether it will be possible to create constructive hypertext in the space of the Web, technically or legally. Technically, most existing Web pages are monolithic constructions, and far from amenable to what Ted Nelson calls “transclusion”—which would be necessary for appropriately including portions of old-style Web pages in a newly constructive Web. However, the World Wide Web consortium is at work on some elegant solutions for deep addressing in monolithic Web pages, Nelson is currently exploring the transclusion possibilities opened by these new standards and other approaches, and at the very

Joyce writes in the context of education, which is fitting not only because Joyce is an educator, but also because the constructive/exploratory distinction can also help describe the differences between Seymour Papert’s LOGO and Computer-Aided Instruction systems (¶28), or between Alan Kay and Adele Goldberg’s Dynabook (¶26) developed at PARC’s Learning Research Group and MacOS or Windows.

Eastgate remains the most prominent commercial publisher of electronic literature, with an early catalog that includes fictions by Joyce, Stuart Moulthrop (¶48), and other leading new media figures. Some of Eastgate’s groundbreaking titles anticipated now-widespread Web literary practices—an example is Deena Larsen’s *Marble Springs*, which invited contributions to its network from readers, and managed to get them, despite the fact that delivery required significantly more effort than hitting a Web form’s “submit” button.

Nelson discusses this using terms such as “windowing”—rather than “transclusion”—in his first version of *Literary Machines* (¶30).

02

08

11

17

21

30

613

47

Video documentation of Kendall's *The Clue* appears on the CD.

least many images are stored in a manner that would allow them to be easily transcluded. Given this, the next matters to consider are legal ones, where the picture is not rosy. A number of lawsuits have already been waged to prevent linking to, and contextualizing of, Web information in a way that its owner considered inappropriate. As the Web continues to integrate with the larger media industry, legal territorialism seems likely to grow only more vicious. Constructive hypertext may be simply be incompatible with our current cultural moment. While some might exclaim that our culture will have to change in the face of hypertext's manifest destiny toward the constructive, evidence for such assertions is lacking. Rather it seems possible that constructive islands will have to exist separately, and then "export" to the Web—or, if constructive hypertext is made part of the "live" Web, that budding builders will have to be very careful of what they include in their reading/writing process.

Meanwhile, just as Seymour Papert's constructionist educational philosophies can be applied with pencil and paper, educators and writers are working to employ constructive hypertext models in the barely-exploratory environment of the Web; Robert Kendall's Connection Muse is one example.

—NWF

CritSuite (from the Foresight Institute) and Annotation Engine (from the Berkman Center for Internet & Society, Harvard Law School) are currently-active, open source projects seeking to bring annotation to the Web. Annotea is a project of the World Wide Web consortium in this area, which uses and helps to advance open Web standards. <<http://crit.org/>> <<http://cyber.law.harvard.edu/projects/annotate.html>> <<http://www.w3c.org/2001/Annotea/>>

Further Reading

Eastgate Systems <<http://www.eastgate.com>>

Goose, Stuart, Wendy Hall, and Siegfried Reich. "Microcosm TNG: A Framework for Distributed Open Hypermedia" IEEE MultiMedia 7:3, Los Alamitos, Calif.: IEEE, 52–60, July–September 2000. <<http://www.computer.org/dsonline/articles/ds3rei.htm>>

Joyce, Michael. *Othermindedness: the Emergence of Network Culture*. Ann Arbor: University of Michigan Press, 2000.

Kendall, Robert, and Jean-Hugues Réty. *Word Circuits Connection Muse*. Web site offering download of the Connection Muse tools. <<http://www.wordcircuits.com/connect/>>

Original Publication
Academic Computing 3 (1988): 10–14, 37–42. Reprinted in Michael Joyce, *Of Two Minds: Hypertext Pedagogy and Poetics* (Ann Arbor: University of Michigan Press, 1995).

Siren Shapes Exploratory and Constructive Hypertexts

614

Michael Joyce

Hypertext and hypermedia are increasingly perceived as instances of a cardinal technology, i.e., tools for working at traditional tasks which have the effect of changing the tasks themselves. Yet, a fairly common reaction to hypertext and hypermedia systems in product reviews, in technical literature, and among everyday users of these tools is the one expressed by Jeffrey Conklin in his still definitive article:

"One must work in current hypertext environments for a while for the collection of features to coalesce into a useful tool." (Conklin, 1987)

This is a kind way to say that you have to figure out what to do with these things. I have spent much of the last four years figuring out exactly that. As a co-developer (with Jay Bolter and John B. Smith) of Storyspace, a hypertext program for the Macintosh computer to be published by Brøderbund Software this year, I have approached what to do with these things as a design question. As a fiction writer seeking to work in a new medium, I have approached it as an artistic question; and as a teacher, a practical, pedagogic, and a sometime political one.

My colleagues and I have had three years experience using hypertext tools in a variety of settings within a comprehensive community college. Storyspace has been used not only as a day to day writing and thinking tool, but also in Chemistry, Nursing, Technical Writing, Creative Writing, Literature, and Developmental Reading and Writing courses.

Users include faculty, community professionals, creative writers, and a wide range of traditional and non-traditional students, from high school age to senior citizens.

Our experience in using Storyspace predates the release of both Guide and HyperCard and began with fairly unstable developmental versions and continues as Storyspace is shaped for commercial release. It has helped us test our notions about learning and informed our work in designing hypertext learning tools using Storyspace and HyperCard both. To some extent, then, given this kind of applied design and testing, I have also approached the question of what to do with hypertexts as a research question.

Given all these claimed approaches, it would seem I ought to have gotten somewhere. I will try to say where that might be, and I will even take the obligatory stab at explaining what hypertext is, as well as what it might become. (Hereafter, I will use the term hypertext where hypermedia would do as well, since nearly all hypertext systems involve other media, and I know of no hypermedia systems which use no text.) Before that, however, I want to discuss—briefly and, in the hypertext spirit, idiosyncratically—some design issues, artistic issues, and ultimately practical, pedagogic, and even political issues which we ignore at our peril. In doing so, I want to offer a somewhat polemical (and, for a hypertext developer, arguably self-serving) description of a set of these perils as well as a set of the promises which accompany them.

Indeed, hypertext tools offer the promise of adapting themselves to fundamental cognitive skills which experts routinely, subtly, and self-consciously apply in accomplishing intellectual tasks. Moreover, hypertext tools promise to unlock these skills for novice learners and to empower and enfranchise their learning. Ironically, however, our ability to deliver upon these promises may be imperiled in the short run by many of the same factors which make this technology so promising.

For instance, the ready adaptability of these tools to more traditional uses is especially compelling given the technological frosting they so easily spread upon stale cake. This disincentive to change is in no way novel, either in the long history of cardinal technologies, or, especially, in the short history of microcomputers in education. The adaptability of HyperCard, for instance, makes it easy to “author” educational software which merely redistributes the command lines of the worst kinds of supposedly interactive,

“drills-and-skills,” CAI software into gaily embossed buttons and peek-a-boo card fields. Like the Applesoft BASIC revolution in educational software which preceded it and which it so clearly resembles, the HyperCard revolution requires us to rely upon skeptical eyes, keep a shrewd ear open to word-of-mouth (or word-of-network) advice, and exercise a cool hand (and fast delete finger) in choosing among a burgeoning list of titles.

Because the price is right (again, like the Applesoft revolution, much of this software is shareware or relatively inexpensive), it is likely that the potential benefits outweigh nearly all the short run perils save perhaps the most crucial one. The peril of over-promising threatens not just to sap the resilience of educators who must wade through the dross and justify the costs. It also threatens the credibility and creativity of innovators who find themselves having to disaffiliate and differentiate before they can discover. Avoiding the peril of over-promising, I will argue below, depends upon our ability to distinguish between what I call Exploratory and Constructive uses of hypertext as a learning tool and our willingness to pursue and encourage the development of both.

By exploratory use I mean to describe the increasingly familiar use of hypertext as a delivery or presentational technology, i.e., as Guide and HyperCard are currently most often used. Exploratory hypertexts encourage and enable an audience (users or readers are inadequate terms here) to control the transformation of a body of information to meet its needs and interests. This transformation should include a capability to create, change, and recover particular encounters with the body of knowledge, maintaining these encounters as versions of the material, i.e., trails, paths, webs, notebooks, etc.

The hypertext audience should also be able to readily understand the elements which make up a particular body of knowledge, plot the progress through these elements, and locate them at will. These so-called navigational capabilities should be present both within the organizational structure of the hypertext and from the perspective of the particular versions of it which the audience creates. At least in the short run, and especially in educational hypertexts, the audience should be able to view alternative visual representations of the corpus's structure and, in some sense, be able to differentiate the unique organizational schemes of hypertext from the more conventional organizations of print and other

media. Ideally, an exploratory hypertext should enable its audience members to view and test alternative organizational structures of their own and, perhaps, compare their own structures of thought with hypertext and traditional ones.

By constructive use I mean to describe a much less familiar use of hypertext as an invention or analytic tool, i.e., the uses we have designed for and made of Storyspace. These are also the uses to which outline processors and their offspring, Personal Information Managers (PIMs) such as Agenda and Grandview, have been put. For that matter, these are also the uses which have sometimes been forced upon word processors, spreadsheets, and databases. Just as exploratory hypertexts are designed for audiences, constructive hypertexts are designed for what Jane Yellowlees Douglas has termed "scriptor[s]." (Douglas, 1987) Scriptors use constructive hypertexts to develop a body of information which they map according to their needs, their interests, and the transformations they discover as they invent, gather, and act upon that information. More so than with exploratory hypertexts, constructive hypertexts require a capability to act: to create, to change, and to recover particular encounters within the developing body of knowledge. These encounters, like those in exploratory hypertexts, are maintained as versions, i.e., trails, paths, webs, notebooks, etc.; but they are versions of what they are becoming, a structure for what does not yet exist. Constructive hypertexts, unlike exploratory ones, require visual representations of the knowledge they develop. They are, in Jay Bolter's phrase, topographic writing.

Like the audience of exploratory hypertexts, scriptors should be able to readily understand the elements which make up a developing body of knowledge, plot their progress through these elements, chart new ones, and locate them at will. These so-called navigational capabilities should be present within both the developing organizational structure of the hypertext and within the particular versions of it which the scriptor discovers. Scriptors must be able not merely to view but also to manipulate alternative visual representations of the invented structure of the hypertext, switching between them with a minimum of intellectual effort, or what Conklin has termed cognitive overhead. They must be able to differentiate and label unique organizational schemes of hypertext, plotting them against more conventional organizations of print and other media, and

generating either kind of organization at will. A constructive hypertext should be a tool for inventing, discovering, viewing, and testing multiple, alternative, organizational structures, as well as a tool for comparing these structures of thought with more traditional ones, and transforming the one into the other.

Transformation of knowledge, I would suggest, is the litmus test we should use in judging both exploratory and constructive hypertexts. It is a critical test in judging whether courseware authored with hypertext tools engages learners in looking at material in new ways, or merely looks like a new way of learning. In many ways, of course, this kind of test is not new to us. Understanding, plotting, navigating, and recreating knowledge structures is the essence of learning. As the current critical thinking across the curriculum craze attests; however, we are less and less certain of our ability to convey these skills.

This uncertainty, coupled with the novelty and ease of authoring with hypertext tools, raises the peril of over-promising. Too often, despite our inbred intellectual skepticism and knowledge about free lunches, we in education have approached technology with what might be called a hunger for automaticity. Both well and poorly designed exploratory hypertexts feed this hunger. It is easy to think that, because learners can move through a body of knowledge in new ways, they know where they are going. We long for a learning machine and think that the computer will do, even as we know that the computer did not do the first time around (and the SRA carrel with its "programmed learning" buttons did not do before it).

Poorly designed exploratory hypertexts often involve a second or third coming of the learning machine. They assuage the hunger for automaticity with the full-bellied inertia of tradition. We know the bulk of this stuff; we have chewed it over for years. It is easy to think that, because learners can move through a body of knowledge in new ways, we know where they are going.

Yet, neither we nor they know as much about learning as either party would think. Paradoxically, we know too much about learning to make such claims. It is a fair bet to say that our age is at least as likely to be known in the future as the Age of Learning as it is as by the ordained cliché, the Information Age. The body of knowledge about learning in psychology, cognitive science, neurophysiology, artificial intelligence, and so on, would itself make a rich exploratory

hypertext. The individual versions we might create as we wend our way through this corpus would likely share certain sets of contrary attributes. Learning is multiple yet integrative, difficult yet universal, not easily schematized yet apparently systematic, inherently personal and yet socially manifested, and so on.

These contraries provide cautionary measures against which to judge exploratory hypertexts as learning tools. Also, they introduce and outline the promise of what I have termed constructive hypertexts. Every well-designed, exploratory hypertext proceeds from a constructive hypertext created by its author or team of authors. The transformation of knowledge which an audience works upon an exploratory hypertext, in some important sense, parallels and rehearses the prior constructive encounter of its authors' associative thought processes. In this instance, however, the word parallel is almost certainly metaphoric language and inadequate at that for a process which is sometimes orthogonal, sometimes congruent, sometimes isomorphic, but always, in some important sense, anticipatory. The authors and audience of hypertexts share a transforming interrelationship. They are, to use an overused term, co-learners. Even the most transparent exploratory hypertexts involve a shared process of mapping this interrelationship, while constructive hypertexts make the transparent mapping visible, active and personal.

The importance of the *process* of associative thinking is suggested in a preliminary evaluation of the Intermedia project at Brown University. Discussing Professor George Landow's groundbreaking work in designing the Intermedia English 32 course, IRIS investigators reported an "unintended consequence of [their] research, discovered when Professor Landow was forced to teach . . . before the [Intermedia] workstations . . . were ready . . . was that he changed the way he organized his course. As a result he felt that students grasped pluralistic reasoning styles far better than in previous years." (Beeman et al., 1987) The authors go on to note that "students were also far more satisfied with the course than in previous years," and present data which show interesting shifts in students' evaluation of both the amount learned in the class and their overall evaluation of the course.

Both sets of evaluations rose significantly when the course was offered before the workstations and exploratory hypertexts were ready for students. When the course was

offered with the workstations, both sets largely maintained gains but, interestingly, fell off to previous levels (or beyond) among students who rated their learning as less—or the class worse—than other classes at Brown. Meanwhile, high end ratings fell off enough to make them of interest.

At least one implication of these shifts is worthy of further research and, more to the point, suggests the potential for constructive hypertexts as instruments of learning. Landow's reorganization of his course might be said to have mirrored his associative (or pluralistic) thought processes in creating a constructive hypertext, i.e., the design for the exploratory hypertext, English 32. As a distinguished scholar and critic, he certainly possessed these associative, pluralistic thought processes well before he set out to represent them in a hypertext. Yet, the benefits of doing so were dramatically perceptible to his students. It would seem to follow that these same benefits ought to be extended to learners themselves, especially if further investigation of such "unintended consequence[s]" yields evidence that exploratory hypertexts yield benefits to instruction which are short-lived or, at least, subject to degradation.

It would only be stirring the hunger for automaticity to suppose that, simply by the instrument of creating constructive hypertexts, students could match the prowess of a practiced scholar in associative thought. However, it would be equally unwise, and something over-promising, to suggest that students could gain that prowess by simply exploring the scholar's representation of it. The Intermedia project, to be sure, does not make either mistake; while largely a vehicle for exploratory hypertexts, it provides powerful constructive tools for learners to use in transforming bodies of knowledge.

The importance of an anticipatory, transforming interrelationship among co-learners may perhaps not represent a novel contribution to our understanding of learning; but hypertext, as a cardinal technology, does offer a novel environment for enabling and exploiting that interrelationship. Constructive hypertexts renew an ancient promise, one which would make us know ourselves and become authors of our learning.

This author's role in this transforming interrelationship is like that of Jane Yellowlees Douglas's "scriptor for the potential experiences of . . . readers." While Douglas concerns herself with hypertext interactive fiction, her insights into the importance of authorial intention hold true for more expository, exploratory hypertexts in which scriptor-learners

intentionally become their own readers. “The yields we select, the defaults we discover, influence our understanding of the contents of the text we read,” says Douglas, “in most cases, we realize that we have, somewhat unwittingly, made certain interpretive or navigational decisions based upon our apprehension of authorial intention.” When the author is oneself, apprehending authorial intention becomes a discovery of one’s own distinctive structures of thought.

Douglas focuses upon the literary-critical implications of this apparent and awkward resurrection of authorial intention as a “subject” of literary texts. Other hypertext theorists such as Diane Pelkus Balestri, Jay Bolter, and Frank Halasz address the relationship between authorial intention and structures of thought more directly. In the February 1988 issue of this magazine, Balestri discusses the “constructability” of what she calls softcopy (echoing a term which, Ted Nelson informs me, he coined in the mid-60s). In Balestri’s usage, softcopy, i.e., text on the screen rather than in print, leads to an understanding of “text as having patterns, often multiple patterns for a single text; [and] defines coherence in terms of linkages among parts of a text.” “Hypertext,” she suggests later, “unlike softcopy, changes the relationship between writer and reader. The reader becomes a collaborator, constructing and reconstructing the text, choosing his own path through it.”

The differences Balestri sees here, between the coherent patterns of links among parts of a text and the constructed patterns which a reader makes, might prove less a distinction between softcopy and hypertext than another description of the interrelationship between sriptor and audience. In any case, Balestri points to the need for training hypertext audiences in the new habits of thought necessary to perceive coherence in patterns and links, and to generate coherent patterns and links of their own.

These concerns are, as I have noted, not much different from the concerns we bundle under the rubric of critical thinking, or the general category of learning. Constructive hypertexts address these concerns in a more conscious way than exploratory hypertexts. They enable audiences of expert and novice readers alike to act as sriptors and focus upon the discovery of coherent structures and linkages; and, most importantly, to use a full range of cognitive skills, especially visual ones, to discover new structures and linkages. Balestri’s notion of softcopy invites us to consider how coherence can be (and is) both visually represented and perceived, and to

consider how we can train learners both to recognize and to generate visual representations of patterns of structure.

My collaborator, Jay Bolter, in his book *Writing Space* argues that these coherent patterns and links are elemental aspects of the associative nature of writing:

... no text is only a hierarchy of elements. A hierarchy is always an attempt to impose rigid order upon verbal ideas that are always prone to subvert that order. The principle of hierarchy in writing is always in conflict with the principle of association. One word echoes another, one sentence or paragraph recalls many others earlier in the text and looks forward to still others. . . . Associative relationships define alternative organizations that lie beneath the order of pages and chapters. . . . Previous technologies of writing, which could not easily accommodate such alternatives, tended to ignore them. The printed book has made the best effort to accommodate both hierarchy and association. . . . The table of contents, listing chapters and sometimes sections, reveals the hierarchy of a text, while the indices record associative lines of thought that permeate the text. . . . An index defines other books that could be constructed from the original book . . . and so invites the reader to read the book in alternative ways.

Bolter proposes that electronic writing, such as hypertext, is an instance of what he calls “topographic writing . . . both a visual and a verbal description . . . not the writing of a place, but rather writing of or with places, spatially realized topics.” Topographical writing is a spatial, visual medium as well as a verbal one. (Another reason I prefer the term hypertext to hypermedia is precisely because hypertext treats *everything* as topographical text; hypertext is the word’s revenge on TV.)

“Although the computer is not necessary for topographical writing,” Bolter notes, “it is only in the computer that the mode becomes a natural and, therefore, also a conventional, way to write.” For the computer “provides a writing surface with an extension and structure unlike previous technologies,” one in which “topics . . . have both an intrinsic and extrinsic significance . . . they have a meaning that may be explained in words, and they have meaning as elements in a larger structure of verbal gestures.”

Frank Halasz focuses upon the distinctive interrelationship between scripted and discovered patterns of structure. Halasz notes that hypermedia systems, of course, require the capability to search for specific content,

i.e., words, keywords, and so on. However, noting that “content search basically ignores the structure of a hypermedia network,” he calls for the development of “structure searches” in hypermedia systems. As Halasz describes them, structure searches are not merely ways of seeking patterned coherence, but in fact ways to identify what Bolter calls elements in a larger structure of verbal gestures. (Halasz, 1987)

As an example of a “complicated structure query, involving an indefinite sequence of links,” Halasz proposes a formulation which, though somewhat of a technical sounding riddle, is nonetheless quite visual: “a circular structure containing a card that is linked to itself via an unbroken sequence of ‘supports’ links.” With a little study this verbal formulation discloses the structure of the riddle. “This query,” says Halasz, “could be used, for example, to find circular arguments.” More importantly, as topographic writing in a hypertext, the visual riddle might likely be solved more easily than the verbal riddle of the query. In this case, the authorial intention, or inattention, of a scriptor would disclose itself in a conscious search for patterned links; and scripted links would become discovered links.

As another of his seven issues, Halasz proposes the need for “virtual or dynamically-determined structures” as a way to eliminate the “static nature of hypermedia networks.” Again, the verbal riddle is more foreboding than the visual one here. Hypermedia networks are static because they only contain patterns and linkages which you, or someone, put there on purpose; they are what Halasz calls “extensionally defined” because the “exact identity of their components” are specified. For example, you might build a structure of everything your student, Betsy, wrote to you while you were away at a conference, but that structure isn’t there until you decide to put it there. You cannot, in fact, see it.

A dynamically-determined structure would just show up whenever you wanted it to. You would, in fact, have to query the hypertext at least once, building what Halasz calls an “intensionally defined” structure. To paraphrase Halasz’s example, you would specify a sub-network containing all the nodes created by Betsy in the last week. However, once that structure was created, every time you looked at it you would see what Betsy wrote during the last week. What you look for would be there to see when you looked. (This notion is not unlike what Alan Kay, the distinguished Apple Fellow, has termed a “software agent.”)

It is not too far-fetched, I think, to suggest that the ideas of Douglas, Balestri, Bolter, and Halasz open the possibility that hypertext learning tools may result in the discovery of what might, in a bad pun, be termed missing links, i.e., novel structures of thought and new rhetorical forms. What may seem far-fetched now is, I think, not just likely but certain. These new forms promise (or, if you will, threaten) to rival, or even supplant, the structures we have come to believe are more god-given than Gutenbergian.

This raises another peril, one perhaps more serious than over-promising. It does take some time to get used to and use from, hypertext tools, especially when you use them constructively. Once you are used to them, they seem to adapt themselves so effortlessly to quite familiar, almost fleeting, and seemingly routine habits of mind that you are hard put to characterize, let alone schematize, them. Because of their fleeting quality, it is tempting to trivialize such habits of mind; because of the difficulty in schematizing them, it is tempting to presume them inferior to more established habits and structures.

To the extent that hypertext challenges traditional intellectual structures, it may be that this cardinal technology, like others before it, will threaten too much to unhinge us. We may perhaps, in the short term at least, lack the vision to appropriate these tools to the new tasks they suggest.

A true test of a challenge to traditional intellectual structures might be whether it is embraced as commonplace by those who do not feel heirs to the tradition. For the past few years, my colleagues Cherry Conrad and Mark Harris and I have used Storyspace as a constructive hypertext tool with students who are often forgotten heirs of a passing tradition. Developmental readers and writers are students assessed as needing further work in writing and reading comprehension in order to succeed in college work. Unlike the technical or creative writing students with whom we have used Storyspace, these students have very little experience with computer programs, keyboarding, or writing environments; and usually possess little or no conscious awareness of formal organizational structures for writing. Compared to these other students, developmental students make minimal use of the complex linking and on-screen hypertext capabilities of our program.

Nonetheless, I want to concentrate on my experience with developmental students including the particular experiences of one student. I offer this account merely as what it is: not

qualitative research or protocol analysis, not research or analysis by any means, but rather that most unstable currency, the teaching anecdote. By doing so, I hope to suggest that the challenge of hypertext to traditional structures can take on commonplace dimensions; and that disenfranchised students, like expert learners, can use such tools to empower themselves in transforming knowledge to their own ends.

Teaching anecdotes of this kind tend to be mini-Odysseys, accounts of where we went, what we saw, and, ultimately, how the world had changed upon our return. To have any sort of mythic power, however, these accounts require some measure of the gates through which we sailed, and how we saw the world as we left. I admit to steering a course between cautionary and visionary pillars. It is cautionary and important in a realm of over-promise to keep in mind Guide inventor P. J. Brown's injunction that "those of us who expect the whole world to rewrite its documentation to fit the needs of our new hypertext system are unlikely to have our expectations fulfilled." (Brown, 1987) This is to say that in using Storyspace with these students, as much as at any time I have taught developmental writing before, I set sail determined, again quoting Brown, to "fit the world as it is, rather than the world as we would like it to be." Developmental students demand as much, although doing so in the most quiet and effective way possible, by ignoring without indication or complaint anything which their very practical experience tells them is bogus.

On the other hand, as teacher, writer, and software designer, I could not avoid navigating by the perhaps visionary truth of what Michael Heim calls "not at all extravagant" assumptions about the future: "Writing will increasingly be freed from the constraints of paper-print technology . . . and vast amounts of information . . . will be accessible immediately below the electronic surface of a piece of writing." (Heim, 1987)

What I saw in my student Les was an ability to see himself as freed of these constraints; constraints which he and his fellow students had, admittedly, freed themselves of long before either by rejecting them outright, or simply failing to learn. What had changed in Les and other students, however, was their ability to perceive and express, as easily as Heim does here, the existence of information below the surface of a writing, and to use that awareness of structure in commonplace fashion to empower themselves.

This is obviously a long way to sail and a fantastical vision to claim to have seen. I should say how we got there. Like many others, the Developmental Education Department at our community college assesses entering students such as Les in the areas of writing, reading comprehension, and mathematics. The writing assessment is a holistically-scored, timed-writing measure in which students write in response to a stimulus. Like most students Les's placement in English 101, Introduction to Writing, was based primarily on his difficulties with organizing a piece of writing and supporting his ideas.

Students are assessed again upon completion of the developmental writing course, and, even before I began using Storyspace with them, my students already demonstrated a high degree of success in post term assessments. I mention this both to avoid any suggestion of automaticity and to suggest another, obvious cautionary note. That is, to the extent that the design of our hypertext writing tool was shaped by my understanding of the writing process, my students might be expected to succeed in using that tool.

Even so, to counterpose a visionary note, success, too, has its variations, few of which can be expected or explained as adaptive behavior. Les succeeded in unexpected and very simple ways, ones which I did not recognize at first and can only incompletely represent now. Quite unselfconsciously and routinely, this eighteen-year old, would-be auto mechanic grasped what was accessible under the surface of an electronic (or, for that matter, traditional) writing and made it his own. Moreover, he did so in a way I had not seen my students do so before, and which, in some sense, made me feel a generational alien, an outdated user of a tool I had helped make.

Mine is what might be called a neo-traditional process model for teaching writing. That is, like many writing teachers I reject pre and post writing distinctions in favor of what is in my case a recursive three-stage model of form-finding, focusing, and shaping. In the form-finding phase, I encourage students to develop an awareness of "impulses toward form" (plans) and "transitory strategies for organization" (goals); using "interrupted automatic writings" to encourage a physiological awareness of shifts in intentionality (impulses or plans) and an ability to visualize and map them graphically on a handwritten page.

I likewise encourage mapping of the consequent shifts in intentionality into provisional organizational structures

(strategies or goals). During initial weeks I use a progressively disclosed hierarchical sequence of abstraction—in which students are requested to consider objects, persons, and ideas in that order—to introduce the concept of “writing toward,” an intentionally fuzzy scheme of organizational approximation.

My students are encouraged to map overlapping strategies (or goals) and to identify changes in impulses (or plans) at their intersections by drawing boxes or frames on their writings. They are likewise encouraged, I should note, to map their own emerging understanding of the writing process against my language for this process, even the most jargony quasi-cognitive-scientific, which I use with them in the same way I do here.

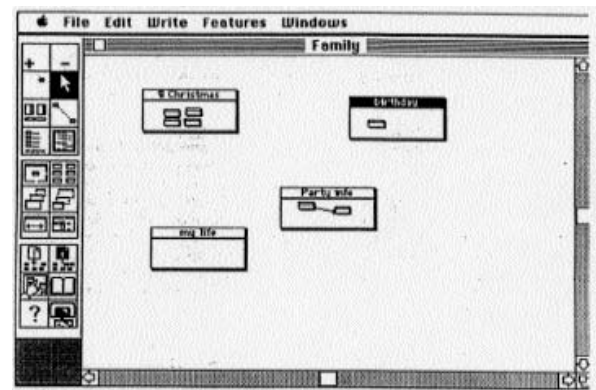
On one level the language is our shared joke against the English teacher world; on another, it is as mythic and empowering as any other part of this Odyssey, a set of more-or-less understood names for skills they possess and have learned to recognize. In any event, this may explain why they are fairly undaunted when they come to use a somewhat complex computer writing environment—something most of my composition colleagues are curious, or even skeptical, about. Most developmental students have long suspected that something familiar lies behind the codes used against them and have developed remarkably complex codes to defend themselves with. Thus, they indulge me in my codes, especially when it comes to computers, which they view, rightly, as the ultimate code machines.

In the focusing phase, students use their cognitive mapping skills to help identify a network of interconnected, though not necessarily sequential, plans which lead to the development of the sequence of goals which will be presented to the reader, thus introducing a recursive notion of shaping. Students are encouraged to continue form-finding recursively whenever sequences are not clear to them, or might not be clear to an eventual reader, thus building a matching, recursive notion of form-finding.

By the time Les came into my course, I had begun introducing Storyspace to my students midway in the consideration of the second stage of the writing process. I showed Les and his classmates how to “port” the writings they had begun during the semester as well as their mapping skills to Storyspace. I encouraged them to use the program as a first stage of the shaping phase, developing an existing writing for

an audience—in this case in-class workshops, and then I, more or less, let them be.

By the time Storyspace came to Les, it had gained some but not all of the features which it will have in release early in the new year. A set of tool icons is arrayed in the familiar Macintosh palette (see Figure 42.1) to the left of a scrollable, empty window in which the structure of a writing and its links are eventually created. The first set of tools let you add and delete the fundamental building blocks of our program, what we call places, in the structure window. Places are individual, editable elements which can contain writing (including graphics); they may be combined into areas which indicate inclusion or subsumption. An area is a place which contains other places. The primary area is the document itself, which may be thought of as a nearly infinite space within which places may be created, connected, ordered, and referenced.



Other tools in this first set let you create automatically linked places for notes; select the pointer for navigation and other operations; use a powerful interactive mini-database to gather and link places in various documents; link places into paths across hierarchies in one or more documents; and choose among what we call Outline, Chart, or Map views of an emerging document.

A second set of window management tools lets you center places in a window; collapse, expand, or organize window displays; enlarge or reduce window sizes; restore the startup size window; and create a copy of the current window.

The third and most powerful tools let you use a “blockbuster” function to create new places from the text of any existing place; use a “blockmaker” function to gather

texts from the various places into one place; use a “placemark” function to create a path of the places you work with or read; navigate through the displays and texts for your documents in either an editable or read-only fashion; get information about documents, places, or text; and link widespread places into paths across one or more documents.

Finally, various menu commands let you import and export texts from word processors or outline processors; distribute sections of text among places; generate new hierarchies from paths; and so on. Throughout its interface Storyspace operates on something of a principle of redundancy to counteract cognitive overhead. We have tried to design an inertialess structure editor in which creating, linking, and rearranging parts of a document can be accomplished in a number of ways, according to your preference, and almost always visually represented.

Powerful word processors often allow you to adapt command keys and menu items to suit your preferences, but largely require you to work within a standardized environment. While Storyspace allows a good deal of this sort of adaptation, it also lets you tailor the environment you work in. Thus, for instance, the outline, chart, and map views of a document are much less specialized environments for different writing purposes than ways of graphically representing the provisional networks and transitory hierarchies which emerge and disappear as different purposes are discovered and developed. So, too, multiple navigation schemes let you perform both content and structure searches of documents; searching for specific parts or exploring at will according to keywords, place names, paths, and so on. Searches are topographic in Bolter’s sense, according to both visual and verbal descriptions, and accessible through both visual and verbal browsing schemes.

Les, like most of our developmental students, worked initially in the map view, the most visual and topographic, and the one most suited to creating constructive hypertexts. In the Map view, you create and name places as individual boxes, each of which contains editable text and graphics, each of which can be dragged and rearranged in clusters on the screen. This clustering ability, of course, offered an easy transition from the classroom work we had done. Clustering, to quote Michael Heim again, “reminds the writer and thinker of the sense of psychic wholeness in the world of increasingly fragmented texts and automated text manipulation.”

I do not wish to dispute Heim’s contention that “clustering cannot properly be done in the computer interface,” especially the kind of clustering he recounts “on a 57 1/2 inch piece of continuous paper,” and especially since his view echoes the fervent criticism and battle cry of hypermedia theorist and IRIS director, Andries van Dam, i.e., “More screen real estate!” All this notwithstanding, our experience does suggest that the map view offers exactly the kind of “expandable graphic or map of thought discoveries” Heim describes clustering as being, and further offers the “sense of wide open creative freedom combined with . . . peculiarities of connection” which he says “no software outliner could permit.”

With all this as prologue, what Storyspace enabled Les to see may seem modest, if not illusory, at first. This anecdote, when finally told, must be told backwards, since I had little sense of what Les had been doing with Storyspace until he did it, and have no memory of giving him any more instruction than the procedural help I gave others as they prepared their writing for workshopping.

Les’s first paper was an uncommonly good attempt at fairly common writing, an autobiography in the form of an account of cars he had used, owned, been ticketed in, and, mostly, wrecked in traffic accidents. It began with his first “borrowing” of his father’s pickup truck and continued through a “links,” an Escort, another Lynx, and ended with his being given the pickup.

As we discussed the printed paper during the last weeks of the course, in a traditional classroom far from the computers, many of us liked one particular sentence toward the end of the paper: “Boy, if he would have had given me that truck when I had asked, this wouldn’t of ever have happened.” We all agreed the sentence managed to say a lot about maturity and desire. However, one particular paragraph, a description of one of these cars, troubled us. We were not certain which of two cars in succession it referred to, and we puzzled over suggestions for placing it while Les listened in silence, according to the custom in my workshops. When it came his turn to respond to our suggestions, he addressed this one first.

“The box is in the wrong place,” he said, “It will make sense when I move it.”

None of us knew what to make of this, and a few students snickered quietly. Les was insistent and confident, moving on.

“I’m glad you liked the part about my father’s pick-up,” he said, “It was the only arrow.”

Neither the sentence in question nor the paper at large said anything about an arrow. The snickering increased, and I began to feel the kind of Lotus-Eater giddiness workshops sometimes bring as they veer out of control. Something in Les's insistence steadied me. He had lashed himself to a certainty we could not see.

"I'm sorry . . ." I said.

A rosy-fingered dawning came over him, and he grinned. "You don't know what I'm talking about, do you?" he asked. He scribbled quickly in the margins on his copy of the paper and held it up. He had drawn a series of boxes in the left margin.

"On the computer," he said, "That stuff about the car belongs next to the Lynx, and I moved the box but forgot to put it back. It makes sense in the right place."

What he was saying was that the structure of the writing existed electronically in a way that he could access it and make it clear to us, his assembled readers.

"And the arrow is a path," I said.

He was pleased that I finally understood, and drew it too in the margin of the paper.

"Yeah," he said, "the only one. I noticed that it starts and ends with my father's truck and so I put that sentence in."

What he was saying was that the verbal formulation of the paper led, topographically, from the visual representation of it. O brave new world, he was saying.

It is to you, I thought. I had missed my first vision of it.

It may be that Conklin's implicit question which I began with, i.e., what to do with hypertext tools, is the wrong question. Perhaps the better question is how to use these things to do better what you already do well. Certainly the early proponents and true visionaries of hypertext believed as much. "The human mind operates . . . by association," claimed Vannevar Bush in describing his Memex in 1945; and from Bush onward through Douglas Engelbart's Augment and even unto Ted Nelson's Xanadu, the visionaries have insisted that the sometimes slippery and obscure trails of hypertext rest upon an underlying bedrock of natural cognition. With nothing less than democratic zeal, each of this Trinity—the citizen-scientist, the engineer-rationalist, and the provocateur-humanist—builds upon a constitutional belief that habits of mind are naturally associative. They see hypertext trails as leading to a kind of shining electronic village upon a hill—an integrated, personalized, machine-enhanced, universally accessible,

associative, new, yet familiar, world platted upon the patterns of synapses, deeded to each according to her or his needs.

It is a compelling and, potentially, an accurate vision; and it is a vision I share. In the forty three years since Bush's first exposition of this vision, it has attracted a litany of adherents. The IRIS project's Intermedia; Halasz, Moran, and Trigg's Xerox PARC-NoteCards; P. J. Brown's Guide; Bill Atkinson's HyperCard; and indeed our Storyspace are each predicated upon equally democratic intellectual principles.

I wonder if we can hold onto them. Until recently the hypertext community, unlike the Artificial Intelligence (AI) community for instance, has been able to make its case incrementally without having to deliver upon strident claims and excessive hype. Ted Nelson's wonderful books and better talks aren't strident, they're fervent (and anyway he baptized hypertext). My mini-Odyssey and gossamer anecdote above borrow on this fervency, as does most anything written in this area.

Even Apple's marketing of HyperCard is less hype than an example of a conversion experience, a sort of corporate speaking in tongues; largely democratic, albeit accompanied by four-color tabloid testimony and Lucasfilm laying on of hands. Conversion experiences are common in the realm of hypertext. In his talks, even as cautious and rigorous an intellectual historian as Bob Jones (whose article appears elsewhere in this issue) reports being filled with wild surmise on first looking into IRIS's Intermedia.

I wonder what these conversions will cause us to overlook. Hypertext, unlike AI, has until recently enjoyed the safe harbor which relative obscurity brings. Yet, education and technological change both stir winds across safe harbors, and now I wonder whether the claims we may find ourselves making about exploratory and constructive hypertexts alike may not put us prematurely in a whirlwind, not unlike the hurricane which accompanied the AI boomlet of recent years. Already a paper suggests developing "a path analysis to classify prominent paths . . . learners take through a hypertext" and then perhaps using "expert systems . . . to help learners access relative portions or sequences of hypertexts." (Jonassen, 1988)

Once I understood what Les was trying to tell me, I quickly secured his permission to let me investigate his workspace and study and report what I found there. While much of what he had done had disappeared in the process of his doing it, I found enough to think that we really ought to

set up our wind machines, “instrumenting” and “journaling” the kinds of behavior he, in his own relative obscurity, had made commonplace.

It was interesting, for instance, to see that the autobiographical writing had sprung from a full-fledged writing environment. A first impulse appears at the end of a place called “Party Info” in a document he named “Family.” The tale of his cars emerged from a typical account of a teenager’s party, resulting from something of what Heim calls the “compensatory discipline” of “releasing,” but which Donald Murray, Peter Elbow, Ken Macrorie, and many other composition theorists have long had other names for.

It was even more interesting to note that he had at some point framed out an organizational outline for the autobiography, deep within a map of another place, called “Christmas,” which included places named for a sequence of gifts—“new toy,” “new bike,” “new car”—none of which he wrote about later.

These uninhabited places upon a map of thought discoveries seemed to call into question the emerging body of research which suggests, as Christina Haas does, that “writers [plan] significantly less when they [use] word processing” and that “there was less conceptual planning and more sequential planning with word processing” (Haas, 1988) And they bring to mind Endel Tulving’s speculations that “the kind of learning reflected in fragment completion and other similar tasks” might be “subserved by . . . an unknown [memory] system . . . the QM system (QM for question mark).” (Tulving, 1985)

Yet, I wonder if all my and others’ speculation will eventually help students like Les learn to do better what they already do well. And I cannot shake the uneasy, liberating feeling that the most dazzling and revolutionary exploratory hypertexts will be developed only when we come, like Les, to create constructive hypertexts which plumb the underlying topographical depths below a surface of boxes, one arrow, and who knows what other siren shapes of thought.