

Toward an ecology of hypertext annotation

Catherine C. Marshall

Xerox Palo Alto Research Center

3333 Coyote Hill Rd.

Palo Alto, CA 94304, USA

Tel: 1-650-812-4288

E-mail: marshall@parc.xerox.com

ABSTRACT

Annotation is a key way in which hypertexts grow and increase in value. This paper first characterizes annotation according to a set of dimensions to situate a long-term study of a community of annotators. Then, using the results of the study, the paper explores the implications of annotative practice for hypertext concepts and for the development of an ecology of hypertext annotation, in which consensus creates a reading structure from an authorial structure.

KEYWORDS: annotation, study, spatial hypertext, reading-oriented systems, consensus.

INTRODUCTION

Annotation is a fundamental aspect of hypertext. In theory, hypertexts grow and change by way of addition – readers respond to hypertexts with commentary, make new connections and create new pathways, gather and interpret materials, and otherwise promote an accretion of both structure and content. In so doing, they crucially augment an existing body of interrelated materials. The foundational work in our field arises from such an annotative perspective: Bush's Memex machine focuses on annotation through trail blazing [2]; Xanadu takes a transclusive approach in which new hypertext seamlessly assimilates portions of older writings [20]; and Augment emphasizes a capacity for Journal system commentary [8].

Ideally, such annotation increases not only the overall girth of the hypertext, but also its value. In his Hypertext '87 keynote address, van Dam justified the capabilities of the groundbreaking FRESS system by saying:

"The reason I encouraged such annotations was that I remembered that when I was in college with Ted [Nelson], I would always grab the dirtiest copy of a book from the library, rather than the cleanest one, because the dirtiest ones had the most marginalia, which I found helpful." [30]

Despite this early emphasis on annotation, there remains substantial work to be done. Many of our systems and methods are (justifiably) geared toward the initial design of a hypertext, not toward the annotation of existing materials (see, for example, [3]). In essence, as Rosenberg points out, they are writing-oriented systems rather than reading-oriented systems [24]. While the role of the reader-as-navigator is widely acknowledged in Web development efforts, the reader-as-annotator is a far less common emphasis. Web link services such as Microcosm's open information services [4], open hypermedia frameworks such as DeVise [12], Web path services such as Walden's Paths [11], and Web annotation services such as CoNoter [7] take important steps in this direction.

In this paper, I take up the question of annotation as a fundamental activity for hypertext readers. I do this through a study that examines the practices of a large, dynamic community of annotators. First, I situate the study within a set of dimensions that characterize existing work on annotation. Then I describe the study itself. The study is used to explore annotation and hypertext in two ways: first from the standpoint of individual practices; then from a perspective that may help us develop an ecology of annotation, one that takes advantage of individual practices to augment hypertexts and increase their value for future readers.

DIMENSIONS OF ANNOTATION

Annotation covers a broad territory. It has been construed in many ways: as link making, as path building, as commentary, as marking in or around existing text, as a decentering of authority, as a record of reading and interpretation, or as community memory. This range suggests a set of dimensions¹ that reflect the forms annotations take (Are they formal or informal? Are they tacit or explicit?); the functions of annotation from a reader's point of view (the degree to which the reader has become a writer, the kind of reading that the reader is engaged in, and the permanence of the marks); and, finally, the roles of annotations as they are used to communicate with others (Are they published or private? Who is the audience?).

¹It is important to consider each dimension as suggesting a continuum, not a dichotomy.

Formal v. informal annotations. An example of annotation at the most formal end of the spectrum is metadata, specifically metadata that follows structural standards and is assigned values using conventional naming authorities. This level of formality helps ensure interoperability: these annotations are, theoretically, more apt to be interpreted in the same way by different query mechanisms. Toward the informal end of the spectrum we find marginalia of the sort that we write to ourselves as we read a journal article. Notetaking tools like *Dynomite*, which uses heuristics to interpret symbols in the margins of an electronic notebook [29], and interpretive tools like *VIKI*, which uses visual and spatial attributes of nodes to infer hypertextual structure [28], help extend the power offered by more formal representations to common annotative practices.

Explicit v. tacit annotations. Many personal annotations, by their nature, are telegraphic, incomplete, and tacit. A highlighted sentence, a cryptic marginal "No!", an unexplained link, a reading history, or a bookmark all pose interpretive difficulties for anyone other than the original annotator (and the passing of time sometimes erodes that privilege). Readers immersed in a text – be it a hypertext or a paper book – seldom make more explicit than that which is required for the task at hand. On the other hand, annotations intended for others to read tend toward the more explicit end of the spectrum. Hence this dimension is most crucially related to intelligibility.

Annotation as writing v. annotation as reading. Readers don't just read. They commune with their documents. They wander, collect, organize, interpret, mark in, and mark on what they gather. The degree to which these annotations are writings on their own forms a dimension. On one end of the spectrum, we find de Certeau's readers:

"Far from being writers - founders of their own place, heirs of the peasants of earlier ages now working on the soil of language, diggers of wells and builders of houses - readers are travellers; they move across lands belonging to someone else, like nomads poaching their way across fields they did not write, despoiling the wealth of Egypt to enjoy it themselves... Reading takes no measures against the erosion of time (one forgets oneself and also forgets), it does not keep what it acquires, or it does so poorly, and each of the places through which it passes is a repetition of the lost paradise." – Michel de Certeau, quoted in [5].

On the other end of the spectrum, we find a postmodern ideal of a polyvocal hypertext, the text as a participatory medium in which annotators are writers, a continuum and tension that Moulthrop explores in [19]. Taken with the other dimensions, it is clear that this dimension is indeed a continuum, and not a dichotomy. Annotations, in their many forms, frequently bridge between reading and writing.

Sometimes this bridging function is surprising, as for example, the photograph of the sky that a student has tucked

between the pages of a meteorology textbook. It is a non-textual annotation that leaves no mark, is informal, tacit, and is, above all a reader's device. Yet it is surprisingly close to a writing – a significant addition to the text.

Hyperextensive v. extensive v. intensive annotation. In his DL'97 paper on reading and attention, Levy takes up the notion of distinctions among hyperextensive, extensive, and intensive reading. Hyperextensive reading involves the link following, fragmentation, and repetition we associate with hypertext; extensive reading is the sort of reading we associate with day-to-day analytic activities – a broad reading of many documents at a time; and intensive reading is a deep engagement (possibly repeated or ritual) with a single text [15]. The same rhythms and contrasts seem to be true of annotation practices. In hypertext terms, it is the difference between link- or structurally-oriented annotations, in which two or more lexia are involved at a time and the entire hypertext is in the foreground, or within-lexia annotations, in which engagement is primarily with the single lexia and the hypertext is in the background.

Permanent v. transient annotations. Annotations, unlike diamonds, may not be forever. If annotations are indeed reflections of a reader's engagement with the text, their value may only hold for the current traversal through the narrative or hypernarrative. On the other hand, some annotations have been observed to bring value to future readers (including the original annotator) [16]. This tension may be at the root of some of the debate concerning the status of annotations; see for example the debate between Sven Birkirts, Robert Stein, Carolyn Guyer, and Michael Joyce in *Feed* magazine [9].

Published v. private. We all know of circumstances in which annotations are a private form – the nasty note scribbled in the margin of something we are reading that we find irritating. Most of our personal annotations, however, are not strictly private: when we give a book to a colleague, we seldom pause to erase all our notes from the margins. Published annotations are also a common form. Annotated editions of important scholarly works are a good example of published commentary. As hypertext matures, there are examples all along this spectrum, especially when hypertext is used as a vehicle for class discussions over a carefully constructed corpus of related materials [14].

Global v. institutional v. workgroup v. personal. It is in this continuum that we see the various visions of hypertext and differing assessments of the value of annotations. Certainly, Nelson's vision (and later, Berners-Lee et al.'s vision) of hypertext and hypertextual annotation is global; Engelbart's vision extends from workgroups or institutions to communities; and Bush's trails were intended for the benefit of an increasingly fragmented on-line scientific community.

Second generation hypertext tools like Intermedia's InterNote facility [6] and the Prep Editor [21] were intended for collaborative authoring, and the most interesting aspects of the tools – InterNote's warm linking that allowed the text

of the annotation to "slide down" the link to replace its anchoring text, and the Prep Editor's ability to support side-by-side comparison of a whole set of annotations – were related to writing within the workgroup. Many current annotation capabilities are built into systems in ways that render them inherently personal; there is no way for different users to share annotations.

The study I describe in the remainder of the paper explores what are primarily informal, largely tacit, annotations made during the course of intensive reading. They were intended by their authors as transient for the most part, since – as I will explain – they were not retained, nor were they regarded as published commentary. Rather they were private notes, scribbles, and markings that were intended as meaningful only to the reader. To wit, some of them were not serious; they may have not even functioned in the way the annotator planned them to. What is observably true is that annotation is a practice that develops over time. Experience and disciplinary expectations change the marks people make.

I use an analysis of these informal (and, in many cases, not intelligible) personal annotations to explore the ways in which these ubiquitous marks might be useful and effective from an ecological point of view in large scale hypertexts – how they may be used to inform some aspects of hypertext functionality and how they may function to benefit other readers without changing the mission or practices of the original annotator.

A LONG-TERM STUDY OF PERSONAL ANNOTATIONS

How can we explore naturally-occurring personal annotations such as the ones that van Dam describes? Because they are personal (and usually not published), they are not accessible for analysis. Because they are idiosyncratic, it is difficult to identify patterns and regularities in them, especially the kind of regularities that might inform hypertext system design. To examine and compare annotations, it is helpful to find multiple copies of the same text, annotated under similar circumstances.

In late 1996, I began conducting a study of annotations based on the marked-up textbooks that were available in an on-campus bookstore at a major university (about 7000 undergraduate and 7000 graduate students attend this university); the bookstore has a 'buy-back' policy that enables students to sell their used textbooks back to the bookstore after a course is over regardless of the kinds of annotations they have made. This arrangement afforded me access to multiple annotated copies of the same edition, as well as to a community of annotators. The study began by looking at the mechanics and uses of marking in books – in other words, the form annotations take, the functions they perform for the reader, and the value they hold for future readers. The initial results are detailed in [16].

In summary, I found great fluidity in form – students used highlighters, pens, pencils, and other writing implements to record marginalia, make symbolic notations, draw on and

over text, write between the lines, underline, circle, box, and highlight all kinds of elements of books. This is no surprise; readers are enormously creative in their engagement with texts. More surprising were the differences in marking practices according to genre, and kinds of annotations that deviated significantly from the kind of commentary and notes we anticipate in our systems designs. For example, a phenomenon that occurs in particularly dense narratives is that annotation becomes a visible trace of human attention. Some textbooks – philosophy texts, for example – contained page after page of reader-highlighted narrative. These markings did not seem to be interpretive. Yet they clearly were important to the physical act of reading.

Although examination of the form and function of the students' annotations was of potential interest from a system design standpoint, the most interesting outcome of the study was the deep ambiguity in the value of the annotations. Some students took van Dam's approach and sought out textbooks with useful-looking markings, usually in the form of longish marginal notes (complete phrases or sentences). Others shunned the marked-in copies, looking for the elusive pristine used textbooks. One clear finding was that the students had developed very strong ideas about what makes annotations valuable or distracting. What implications does this ambiguity hold for the future of annotations in the Docuverse?

I continued the study in the university bookstore, shifting my emphasis to more speculative aspects of annotation. First, how may annotations on a page inform hypertext concepts? They are visual and spatial records of interpretation; they incorporate all kinds of useful strategies and shorthands; and they exhibit a wide range of linking and anchoring styles. Second, given the actual ways students annotate their texts, what are the ways in which we might implement a more 'ecologically sound' facility for sharing personal annotations, taking seriously concerns about both utility and privacy? While Web tools and hypertext systems before them have emphasized intentionally shared annotations, it seems important to explore how we might use these other kinds of ubiquitous annotations.

The extended study involved four different sorts of data collection: examining the used textbooks one-by-one; observing students as they sorted through the used textbooks and discussed among themselves which one to purchase; interviewing students as they (and I) looked through the used textbooks; and performing detailed comparisons across a specific set of annotated copies.

Textbooks as annotated artifacts. The annotated textbooks are readily available for examination. As in the earlier part of the study, I paged through all (or a large, representative subset) of the annotated copies that were available for a given textbook. All told (including the 150 books used in the earlier study), I examined 410 books representing 39 titles in 21 different subject areas.

Observations of book buyers. The students often came to the bookstore in companionable small groups and discussed their purchases. For the fall semester, freshmen were sometimes accompanied by their parents, with whom they explored their textbook options. Because there were so many used books available, the students (or student/parent pairs) would often help each other choose the "best one" from a stack. Their conversations, and their choices from stacks of books I had already gone through, were a vital source of insight about the value of the annotations.

Informal interviews of book buyers. Informal interviews with students buying used books provided a valuable window onto the students' own annotation practices as well as their assessment of the value of their peers' annotations. Talking directly to the students also proved vital to establishing the ways in which annotation developed as a practice. First-term freshmen begin with a very sketchy idea about how (or whether) they will write in their textbooks. Upperclassmen are far more savvy about their own marking practices and the value of the annotations made by others. This is born out by looking through books used in upper division courses, and comparing them with books used in freshman core courses.

Detailed comparison of annotated copies. While it is easy to discover general patterns of annotation, it is also useful to perform a more detailed comparison of annotated copies of a particular book. A detailed comparison involved a line-by-line analysis of the markings students made in the available copies of selected textbooks.

Taken together, the artifacts, interviews, and observations gave me a tantalizing view onto the practice of annotation in a wide range of textbooks. Given the opportunity, it would also have been helpful to watch the marks as they were being made or to interview students who were selling books back to the bookstore (since that would have given me access to the annotator him or herself).

PAPER BOOKS AND HYPERTEXT

Why look at paper books when what we are concerned with is hypertext? Although certainly it would be folly to become enmired in imitating paper systems and paper-based practices, it is important to look beyond the existing on-line facilities for readers. Annotation on paper is a well-supported practice (witness the variety of highlighters, pens, clips, post-its, and other technologies) and admits no shortage of practitioners.

More crucially, an examination of these paper books is readily convincing: annotations on paper are hypertextual. They exist in non-linear relationships to the printed linear text: they interrupt linear reading, are orthogonal to it, connect disparate passages, and in general function as hypertext is intended to. They are playful², informal, serious,

informative, cryptic, and everything in between; in fact, they are a direct reflection of a reader's engagement with the text. It is this engagement with the text that our systems may seek to promote.

There are on-line Web-based tools that support shared commentary (see, for example, CoNoter [7], the Foresight Institute's CritSuite [10], Phelps's and Wilensky's Multivalent Document annotation facilities [23]), but most of these are not really intended for the lightweight personal annotations one encounters in the varying circumstances that constitute reading. There are also a wide variety of personal annotation tools incorporated in document-oriented systems and help facilities. Yet, with the exception of prototype tools like XLibris [26], it is rare to encounter support for making the kinds of fluid annotations one sees on paper (which is due in part to the awkwardness of marking using a mouse). Hence the writings and markings in paper books can be a viable place to uncover new insights into hypertext.

What can we learn about existing hypertext concepts like associations, anchors, and types from marks on paper? And, given a reader's perspective, how can we strengthen these basic facilities in our systems? Below I discuss how five common hypertextual elements are realized in annotations on paper, illustrated with examples from annotated copies of Plato's *The Symposium*.

Associations, links, and relations. Given a purely visual and spatial medium, how do readers associate annotations with the printed document elements? There are at least four kinds of associations that readers make in the large sample of textbooks. First, they make associations at level that we might consider as 'collection' or composite level. In this study, I observed annotations that refer to many subparts of a single document; for example, a reader makes a note that refers to all of Chapter 7. Second, readers make node-to-annotation links. This is manifested by annotations that don't visibly refer to any particular document element, but are localized within a document part – a longish note written orthogonally to the printing on a page, for example. Third, they make "standard" hypertext associations – from an anchored portion of the text (see the discussion of anchors below) to a note or commentary (see Figure 1). Finally, they make word-to-word associations (in effect pushing the grain-size of the hypertext to the morphemic level [25]); this is particularly common in foreign language texts, in which the student translates a word into his or her native language, usually writing between lines of text.

What is the mechanism for making these associations? Marginal notes are linked to document elements in three quite different ways. The first literally uses arrows to connect an anchored document element to its annotation. The second uses a bracket, brace, or some other mark to associate commentary with text. Finally, readers rely on proximity alone to connect their own marginal (or interline) jottings with the text. These findings are very much in line with the

²For example, an annotation in a C textbook declared to future readers: "See_Ya_later!"

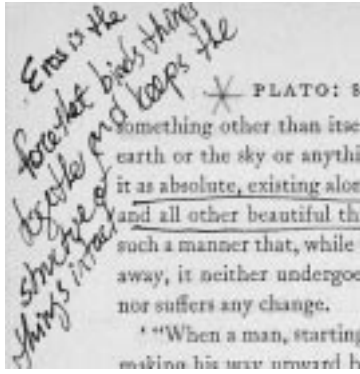


Figure 1. Associating an annotation with a text span.

kinds of implicit associations we saw in our early studies of spatial hypertext [17].

Anchors. Hypertext models have long incorporated a notion of anchors as a way of setting off a span of text, usually as a start- or endpoint for a link. Readers' annotations pose some interesting questions for the anchor construct. Much like associations that remain implicit, the extent of an anchor can be similarly vague. Spanning marks – variations of angle brackets, braces, straight brackets, and lines – may be used to distinguish a region of text. These are loose designators; they do not require the reader to assume the overhead of specifying exactly where an anchor starts or ends (see, for example, Figure 2). Even seemingly well specified anchors – underlines and highlights – either follow existing textual structure (syntactic, like sentences, layout, like boxed-off theorems, or typographical, like italicized terms or concepts) or may be surprisingly quixotic – where they start and end is not a carefully-contemplated decision. In short, the well-structured hypertext of the writer is not the kind of hypertext created by the reader.




Figure 2. A bracket used as an anchor.

Emphasis. Emphasis is not a characteristic we normally associate with hypertext, although the notion of bookmarks and bread crumbs (see [1]) or visual markings in spatial hypertext (see [18]) may come close. Emphasis is the practice of making extra marks next to some (hyper)textual element (usually an anchor, or less commonly, by a handwritten note in the margin) to indicate "This is important." This annotative practice is ubiquitous and significant; it allows the reader to sort through his or her annotations and, by some measure, organize them. Emphasis marks are most commonly stars or asterisks, although the inventive annotator may use a variety of symbols. These

marks may also implement "levels of importance;" for example, two stars may set off a particularly key element.

Figure 3 shows a conventional way that emphasis is implemented, a symbol (or in this case, multiple symbols) associated with anchoring text by adjacency. Other mechanisms for showing emphasis include varying the width of a highlighting mark (from thin to thick) or the use of two different colors of highlighting pen, one layered on top of the other.




Figure 3. Emphasis.

Constructing new nodes from document segments. Some annotations re-segment the document; this often happens when the author's structure does not suit the reader's purposes. How can a reader make new nodes in a paper document (short of tearing out pages or making copies, cutting them up, and reorganizing them)?

Figure 4 shows an example of re-segmentation; the reader has decided to single out several pages (with a line down the side of the visible page, extended onto subsequent pages) and further hierarchically segment the page by numbering passages. Another very common way to re-segment texts is by switching marker colors. This should not be confused with a notion of type; the addition of new structure is heralded by the color switch, not by the color itself.

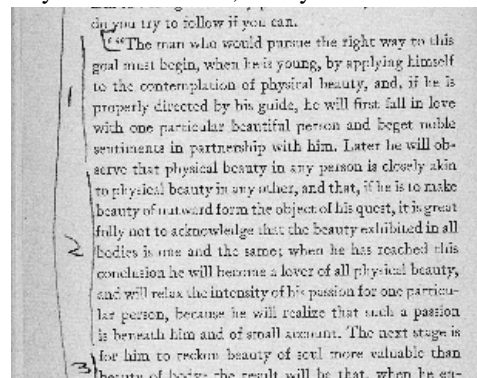


Figure 4. Resegmenting the text.

Types and categories. While most symbol and color use implements either emphasis or re-segmentation, there are notable exceptions. In some cases, the annotations are given types based on color or some other visual property. Figure 5 shows an example of the "key" of a color-based strategy.

We must be very careful, however, not to overvalue the use of even informally typed annotations; color- or symbol-based annotation strategies are less common than the other

phenomena listed above. In fact, in an interview a textbook-buyer confessed to switching among highlighter colors simply to maintain interest in the text.

These hypertextual elements are portrayed in this analysis as static; but if we take to heart the dimension of permanence and transience, it is important to note that these marks are part of a dynamic engagement with a text. The marks alone don't capture their role in re-readings, or in later use of the book. In the next section, I take on the challenge of finding the potential longer-term value of these common types of personal annotations.

AN ECOLOGICAL VIEW OF PERSONAL ANNOTATIONS

Some annotations - annotations of the type van Dam referred to - have clear value to future readers. Interviews and observations of used textbook buyers confirmed that some of the readers' annotations are considered valuable by other readers. But it is also the case that among experienced used textbook buyers, this value is not held as universal across all kinds of markings; experienced used textbook buyers tend to prefer longer written marginal notes over highlightings and text emphasis [16]. According to one book buyer, annotations that look the most like "notes taken in class" - that is, added material with some authority rather than another student's personal interpretation - have the most assured value. This student also considered purchasing annotated books if she knew who the annotator was, and that the annotator was "really smart."

Yet from this study, I found that much of the marking in the books is simply text selection - highlighting, underlining, adding span notations - and emphasis - stars, asterisks, and other symbols in the margin to indicate that a particular text selection is important. Other common markings may serve only to re-segment text or restructure the content. Given the possibilities for marking up electronic text, how might we derive value from the most ubiquitous annotative forms? Is there any way of using these annotations (cryptic jottings, emphasis symbols, underlining and highlighting) in the Docuverse?

Recalling the dimensions introduced earlier in the paper, we also must remember that these are not only personal annotations; they are private annotations that cross inadvertently into a public space. This crossing from private to public is not uncommon in our paper-based document sharing. In many circumstances, people share documents with personal annotations within their workgroups. In this case, privacy is protected mainly by virtue of anonymity³. Any scheme we might consider that involves personal annotations would require careful consideration of issues of privacy and anonymity.




Figure 5. A key to highlighter colors.

If we consider ways of using these annotations, we also must take into account the seamlessness with which the annotations enter public life. No intentional action (beyond selling the books back to the bookstore, an action I assume has a strong financial motivation) is necessary on the part of the annotators; they never need to contemplate whether or not their annotations will benefit other students, nor do they need to do anything extraordinary to release them.

To further investigate how the polysemous annotations and markings might come into play in an ecology of hypertext readers, I performed a detailed analysis of the markings in copies of a particular textbook. This facilitated a very close comparison of the annotations, since the books started out materially the same, and offered identical affordances for note-taking and marking (the exact same form and content).

As a basis for analysis, I used the six available marked-up copies of *Understanding Computers and Cognition* by Winograd and Flores. The textbook is required for an upper-division Computer Science course at the university. *Understanding Computers and Cognition* met some important criteria as the basis for analysis:

First, it is generally used in upper division courses. Since annotative practices develop over time, it was vital to find a text that was marked up by experienced annotators. Texts for lower-division courses had unsustainable and unsustainable annotation strategies; this observation was corroborated by interviews with incoming freshmen, many of whom had never written in their textbooks before.

Second, the book crosses disciplinary boundaries, including philosophy, some biology, and computer science. Although this does not ensure that the results are readily generalized, it does mean that the students will engage with the text as not entirely familiar subject matter, and it will require a close read (and therefore will invite mark-up). As Table 1 shows, annotation extended to much of the textbook. The final chapters of the text (11 and 12) were a notable exception; they were not included in this analysis.

Third, *Understanding Computers and Cognition* is a coherent narrative, rather than a looser collection of subject matter materials. Unlike other upper-division books I examined, many of the readers annotated through large portions of the text, thus making it feasible to perform the labor-intensive analysis on a smaller number of copies.

Finally, the textbook does not contain subject matter that is memorization-intensive; memorization-intensive subject matter tends to result in highlighted terms and definitions

³There is notable evidence in this study, however, that the students did not consider privacy issues when they sold the books back to the bookstore. Names and social security numbers, credit card slips, and other means of identifying previous book owners found their way into the used book stacks.

that, according to my observations, more or less echo the typographical conventions used in the text. Instead, the book contains a complex argument that requires focused attention, the kind of reading we might hope to support in our hypertext systems.

By looking across the marked-up chapters of the book, I hoped to identify n-way consensus, places in the text that all n of the readers (the readers who had marked in a particular chapter) had agreed were important, or at least worthy of pulling out from the text. Was there n-way consensus on what mattered? And if there proved to be consensus, what could a hypertext system developer consider doing with the points of intersection and emphasis? I also hoped to identify other general patterns in the markings to much the same end.

To identify points of n-way consensus, the annotations must first be normalized. They then must be counted and compared. I did this according to the author's division of the work into numbered chapters and sections. That is, I performed the comparison on a per-chapter and per-section basis; this segmentation reflects the readers' tendencies to skip chapters (and possibly sections - certainly there were indications that readers did not approach every section of the book with equal attention and engagement).

Normalizing the annotations. Naturally, every reader has his or her own mode of marking in documents to meet the demands of the situation. All six annotators were fairly consistent - they all used highlighting or underlining, and most indicated emphasis in obvious ways - usually asterisks or stars by the paragraphs. Most wrote some very brief notes adjacent to key portions of the text, usually a recapitulation or summary of what was in underlined or highlighted text.

To normalize the markings, I identified the most common highlighted text unit, which turned out to be sentences (or slight variations on sentences, like sentences with the subject pronoun or noun phrase omitted, or sentences with parentheticals omitted). This normalization tends to eliminate effects introduced by inclusion or omission of articles or subject pronouns, but occasionally tends to obscure differences between picking out a key phrase to remember and emphasizing an explanation of a concept.

I also recorded all instances of extra emphasis that readers marked and the sentences they were associated with. I noted the typographical emphasis of the book as well, since in many of the other books I examined, highlighting occurred in conjunction with typographic emphasis.

There were a number of short phrases written in the annotated chapters of *Understanding Computers and Cognition*. Because most annotators use the implement that is to-hand, many of these are written in highlighting pen; the difficulty of writing with a highlighting pen may have had some bearing on the length of these notes. These notes were used as either labels for particular subsections - additional segmentation - or as emphasis - they pulled out particular

words and phrases from a passage. In the second case, I counted them as emphasis in my analysis.

Annotation counts. Table 1 summarizes the density of annotations on a per chapter basis (percentage of sentences annotated). As is clear from the table, some of the students annotated more vigorously than others. One chapter in copy 2 of the book has an annotation rate of 54%. Other students are far less lavish with their pens; the owner of copy 3 tended to single out a small number of sentences per chapter. We can also see that only one of the annotators, the owner of copy 1, continued to mark throughout the whole book (although a substantial decrease in enthusiasm is observable in the final few chapters). This example illustrates what I found in the larger sample of close to 400 textbooks - annotation is localized, and even the most carefully-designed best-intentioned scheme of personal annotation tends to drop off over time. The experienced used book buyers had much the same observation: I overheard one student telling another to not be put off by annotations early in the book, that "sometimes the writing's only in the first 10 pages or so." This natural drop-off in attention must be considered in any scheme that re-uses annotations.

Consensus analysis. Table 1 shows how often the annotators concurred on their selections. Given the disparity in annotation practices, it seems natural to wonder whether the agreement is simply random - if several annotators marked a quarter or a half of the text, it is likely they will have marked at least some of the same passages. However, if we calculate how many sentences would overlap strictly based on probability, it becomes evident that the n-way consensus is meaningful. For example, Chapter 2 has been annotated by 3 people. They have annotated 21, 54, and 4 percent of the sentences respectively. If we consider that there are 201 sentences in Chapter 2, we might predict that there would be about (rounding up) 1 sentence in the overlap. There are, in fact, 10. Figure 6 illustrates n-way consensus. Table 1 gives the predicted value in parentheses.

As a check, Table 1 also shows the by-chance predicted number of two-or-more-way annotations (again, in parentheses).⁴ While two-or-more-way consensus is not as far off predicted levels, some of the less significant levels of consensus may still be important, especially if we consider the n-way consensus as having varying influence on the selection of surrounding sentences. Often, the selection extends either to the previous or to subsequent sentences.

Does n-way consensus simply reflect the sentences that the authors elevated in importance (i.e. the opening sentence of a

⁴The probability of two-or-more-way annotation consensus was calculated by finding the probability that any sentence been marked by either exactly zero or 1 annotator; the predicted number of 2-or-more-way overlaps was obtained by subtracting this probability from 1 and multiplying by the number of possible sentences in the chapter.

12. the pervasive importance of shared background has major consequences for the design of computer systems.

Artificial intelligence is an attempt to build a full account of human cognition into a formal system (a computer program). The computer operates with a background only to the extent that the background is articulated and embodied in its programs. But the articulation of the background is a never-ending process. In order to describe our pre-understanding, we must do it in a language and a background that itself reflects a pre-understanding. The effort of articulation is important and useful, but it can never be complete.

This limitation on the possibility of articulation also affects more concrete issues in designing computer tools. If we begin with the implicit or

copy 1

12. the pervasive importance of shared background has major consequences for the design of computer systems.

Artificial intelligence is an attempt to build a full account of human cognition into a formal system (a computer program). The computer operates with a background only to the extent that the background is articulated and embodied in its programs. But the articulation of the background is a never-ending process. In order to describe our pre-understanding, we must do it in a language and a background that itself reflects a pre-understanding. The effort of articulation is important and useful, but it can never be complete.

This limitation on the possibility of articulation also affects more concrete issues in designing computer tools. If we begin with the implicit or

copy 4

occur within a shared background. As we will show in detail in Chapter 12, the pervasive importance of shared background has major consequences for the design of computer systems.

Artificial intelligence is an attempt to build a full account of human cognition into a formal system (a computer program). The computer operates with a background only to the extent that the background is articulated and embodied in its programs. But the articulation of the background is a never-ending process. In order to describe our pre-understanding, we must do it in a language and a background that itself reflects a pre-understanding. The effort of articulation is important and useful, but it can never be complete.

This limitation on the possibility of articulation also affects more concrete issues in designing computer tools. If we begin with the implicit or

copy 5

occur within a shared background. As we will show in detail in Chapter 12, the pervasive importance of shared background has major consequences for the design of computer systems.

Artificial intelligence is an attempt to build a full account of human cognition into a formal system (a computer program). The computer operates with a background only to the extent that the background is articulated and embodied in its programs. But the articulation of the background is a never-ending process. In order to describe our pre-understanding, we must do it in a language and a background that itself reflects a pre-understanding. The effort of articulation is important and useful, but it can never be complete.

This limitation on the possibility of articulation also affects more concrete issues in designing computer tools. If we begin with the implicit or

copy 6

Figure 6. An example of n-way consensus in *Understanding Computers and Cognition*.

section, typographically emphasized words or sentences, or set-apart lists and quotes)? An examination of the 17 sentences that exhibit n-way convergence shows that our annotators converged on important sentences buried within long sections. Only eight out of the 17 were typographically-perspicuous parts of a section, specifically list items. The other nine were important, but inconspicuous, sentences within the sections.

Using consensus. I gathered the intersecting sentences and began to investigate how they might be used, for there

appears to be significant value in this sort of consensus. One way to think about these points of convergence is as transitions from authorial structure to new kinds of structure based on readers' consensus. Figure 7 borrows from Nelson's notion of stretch text [20] to show how structure derived from readers' activities might be realized; the figure uses the results of the analysis I have just described. The lexia on the left side of Part A shows the authors' structure for the book. On the right is an expansion based on one instance of n-way consensus in this section. Part B shows the gradual unveiling of more levels (and non-contiguous portions) of text.

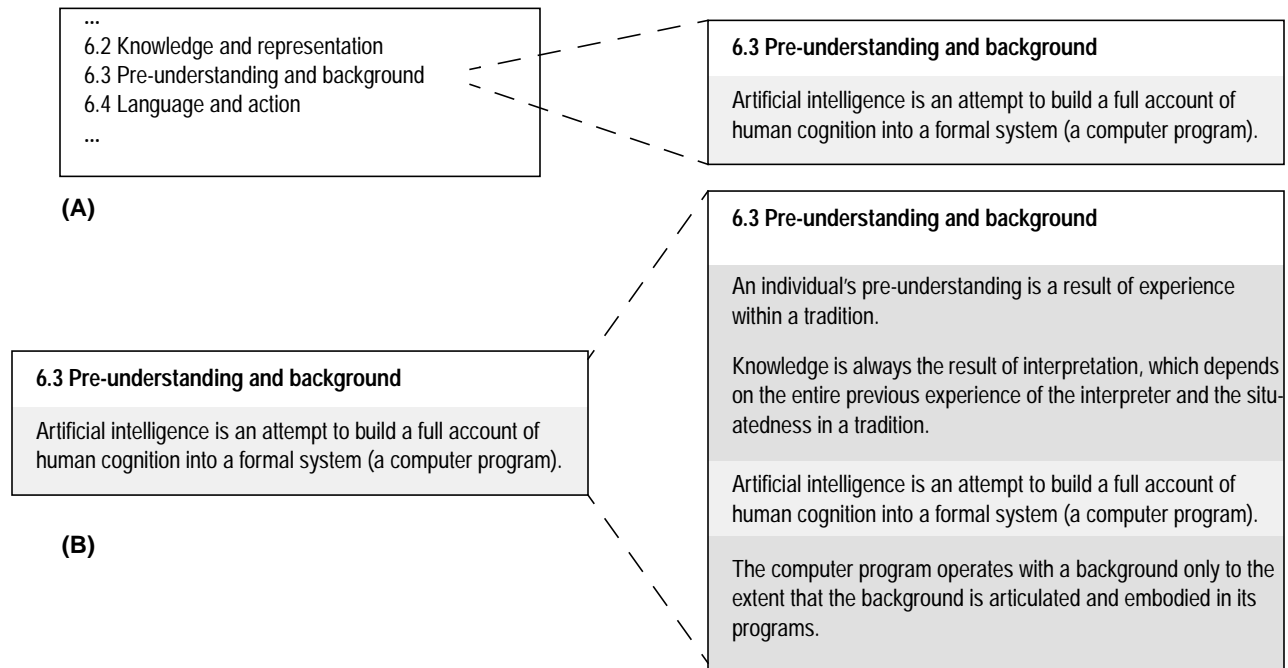


Figure 7. Moving from authorial structure through degrees of reader consensus. Part A shows an expansion from authorial structure to readers' consensus structure. Part B shows an expansion from n-way consensus to n-1 way consensus.

Using a scheme of reader consensus and reader emphasis marks, levels of detail may be unveiled, the text may be re-segmented, or the consensus may be used as a summary (bearing in mind that a common use for annotation is for document review [22]). The mode of display I have suggested in this example fits into spatial hypertext models, and with the fluid link interface described by Zellweger et al. in [32].

Unlike various Web-based community rating services like Ringo [27] or Bellcore’s video rating system [13] that require readers to assign ratings or intentionally evaluate materials for other readers, this type of scheme is a natural outcome of the activity of reading. It is comparable instead to the practices of Webmasters who monitor the link traversal patterns of readers and use these patterns to restructure and tune Web sites [31].

CONCLUSIONS

Is an informal, ecological approach to hypertext annotation feasible? Certainly, personal annotation tools are on the horizon that will enable people to engage with documents and hyperdocuments in a way that aligns with observed practice [26]. So we might expect to see just the sorts of markings I gleaned from the study set of 410 textbooks and the kinds of overlaps I observed in the smaller subset.

However, I have not addressed the problem of gathering these varied personal annotations. If we assume a scheme in which readers make electronic markings on electronic text, it is conceivable that these markings could be gathered and aggregated. Sophisticated annotation servers (for example, see [23]) are in the works. If we assume that the marks are still made on paper – perhaps on printouts of electronic text – then it is necessary to pose a more computationally-intensive scanning-and-analysis route.

Finally, does this kind of consensus-based analysis of reader activity scale? Clearly, full n-way intersections are increasingly rare as n grows larger. But what of lesser degrees of consensus? Even in our small sample, it is evident that there is meaningful consensus at less than full intersections. This is particularly true if we factor in indicators of emphasis, passages that readers have marked as important.

As the number of hypertexts grow, it is increasingly important to find new ways to bring value to them, especially in ways that not only fit with practice, but actively take advantage of it.

ACKNOWLEDGMENTS

I’d like to thank Randy Trigg and Polle Zellweger for their valuable comments and suggestions. Thanks too to Marshall Bern for his insight on how to calculate the probabilities.

REFERENCES

- [1] Bernstein, M. “The Bookmark and the Compass: Orientation Tools for Hypertext Users.” *SIGOIS Bulletin* 9, 1988, pp. 34-45.
- [2] Bush, V. “As We May Think.” *Atlantic Monthly*, Aug. 1945, pp. 101-108.
- [3] *CACM Special Issue on Designing Hypermedia Applications* 38(8), August 1995.
- [4] Carr, L., Hill, G. DeRoure, D., Hall, W. and Davis, H. “Open information services.” In *Computer Networks and ISDN Systems* 28, 1996, pp. 1027-1036.

Chapter	total number of sentences	% of sentences annotated (by copy)						n	n-way consensus actual (predicted)	2+ way consensus actual (predicted)
		copy 1	copy2	copy 3	copy 4	copy 5	copy 6			
1	189	30	21	15	7	9	-	5	- (<<1)	47 (35)
2	201	21	54	-	-	4	-	3	10 (~1)	37 (27)
3	187	11	38	-	7	29	-	4	- (<<1)	54 (36)
4	243	21	16	-	5	11	-	4	1 (<<1)	33 (25)
5	287	13	4	-	4	4	-	4	- (<<1)	14 (5)
6	167	40	-	-	22	13	23	4	2 (<<1)	44 (41)
7	157	29	-	6	-	3	-	3	3 (<<1)	10 (4)
8	242	24	-	10	-	2	10	4	1 (<<1)	29 (17)
9	313	7	-	1	7	-	-	3	- (<<1)	11 (<<1)
10	218	<1	-	6	9	-	-	3	- (<<1)	4 (<<1)

Table 1: Identifying points of consensus in the text

- [5] Chartier, R. *The Order of Books*. Stanford, California: Stanford University Press, 1994.
- [6] Catlin, T., Bush, P., Yankelovich, N. "InterNote: Extending a Hypermedia Framework to Support Annotative Collaboration." *HT '89 Proceedings*. New York: ACM, 1989, pp. 365-378.
- [7] Davis, J.R. and Huttenlocher, D.P. "Shared Annotation for Cooperative Learning" *Proc. of CSCL'95*, pp. 84-88.
- [8] Engelbart, D.C. "Authorship Provisions in Augment." *Proc. of 28th IEEE International Conference*, San Francisco, CA (Feb. 27- Mar. 1, 1984), pp. 465-472.
- [9] *Feed Magazine*. "Page Versus Pixel." <http://www.feedmag.com/95.05dialog1.html>.
- [10] Foresight Institute. "Enhancing the World Wide Web: Social Software for the Evolution of Knowledge." <http://www.foresight.org/WebEnhance/index.html>, 1997.
- [11] Furuta, R., Shipman, F., Marshall, C., Brenner, D., and Hsieh, H. "Hypertext Paths and the World-Wide Web: Experiences with Walden's Paths." *Proc. HT '97*, New York: ACM, pp. 167-176.
- [12] Gronbak, K., Bouvin, N., and Sloth, L. "Designing Dexter-based hypermedia services for the World Wide Web." *Proc. HT'97*, New York: ACM, pp. 146-156.
- [13] Hill, W., Stead, L., Rosenstein, M., and Furnas, G. "Recommending and Evaluating Choices in a Virtual Community of Use." In *Proceedings of CHI'95*, Denver, CO (May 7-11, 1995), pp. 194-201.
- [14] Landow, G.P. *Hypertext: The Convergence of Contemporary Critical Theory and Technology*. Baltimore: Johns Hopkins University Press, 1992.
- [15] Levy, D.M. "I read the news today oh boy: reading and attention in the digital library." in *Proc. DL'97*, Philadelphia, PA (July 23-26, 1997), pp. 202-211.
- [16] Marshall, C. C. "Annotation: from paper books to the digital library." in *Proc DL 97*, Philadelphia, PA (July 23-26, 1997), pp. 131-140.
- [17] Marshall, C. and Shipman, F. "Searching for the Missing Link: Discovering Explicit Structure in Spatial Hypertext." *Proc. HT93*, New York: ACM, pp. 217-230.
- [18] Marshall, C., Shipman, F., and Coombs, J. "VIKI: Spatial hypertext supporting emergent structure." *ECHT '94 proc.*, New York: ACM, pp. 13-23.
- [19] Moulthrop, S. "you say you want a revolution: hypertext and the laws of media." in *Essays in Postmodern Culture*. (Amiran and Unsworth, eds.) New York: Oxford University Press, 1993, pp. 69-97.
- [20] Nelson, T. *Literary Machines*. South Bend, Indiana: The Distributors. Edition 87.1. Also 6th Edition, 1984.
- [21] Neuwirth, C.M., Kaufer, D.S., Chandhok, R., and Morris, J.H. "Issues in the Design of Computer Support for Co-authoring and Commenting." in *CSCW '90 Proc.* New York: ACM, pp. 183-195.
- [22] O'Hara, K., Smith, F., Newman, W., and Sellen, A. "Student Readers' Use of Library Documents: Implications for Digital Library Technologies." to appear in *Proc. of CHI '98*, New York: ACM Press.
- [23] Phelps, T.A. and Wilensky, R. "Multivalent Annotations." in *Proc. of the First European Conference on Research and Advanced Technology for Digital Libraries*, 1-3 September 1997, Pisa, Italy.
- [24] Rosenberg, J. "The Structure of Hypertext Activity." *HT '96 Proc.*, New York: ACM, pp. 22-30.
- [25] Rosenberg, J. "The Interactive Diagram Sentence: Hypertext as a Medium of Thought." *Visible Language*, 30.2, pp. 103-116.
- [26] Schilit, B.N., Golovchinsky, G., and Price, M.N. "Beyond Paper: Supporting Active Reading with Free-form Digital Ink Annotations." to appear in *Proceedings of CHI'98*, New York: ACM Press.
- [27] Shardanand, U. and Maes, P. "Social Information Filtering: Algorithms for Automating 'Word of Mouth.'" In *Proceedings of CHI95*, Denver, Colorado, (May 7-11, 1995), pp. 210-217.
- [28] Shipman, F.M., Marshall, C.C., Moran, T.P. "Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information." *Proc. of CHI '95*, New York: ACM Press, pp. 346-353.
- [29] Wilcox, L.D., Schilit, B.N., and Sawhney, N. "Dynamite: A Dynamically Organized Ink and Audio Notebook." In *Proc. CHI 97*, New York: ACM Press, pp. 186-193.
- [30] van Dam, A. "Hypertext '87 Keynote Address." *CACM* 31(7), July, 1988, pp. 887-895.
- [31] Yan, T.W., Jacobsen, M., Garcia-Molina, H., Dayal, U. "From user access patterns to dynamic hypertext linking." In *Computer Networks and ISDN Systems* 28, 1996, pp. 1007-1014.
- [32] Zellweger, P., Chang, B., and MacKinlay, J. "Fluid Links for Informed and Incremental Link Transitions." In *Proc. HT'98*, Pittsburgh, PA (June 20-24, 1998).