

An Outline for a Functional Taxonomy of Annotation

Allen Renear, Steve DeRose, Elli Mylonas,
Scholarly Technology Group, Brown University
and Andries van Dam
Computer Science Department, Brown University

First Draft Feb 22, 1999
Last Revision Mar 5, 1999
Presented at Microsoft Research, Redmond WA, April 1999.

About this Document

This document presents an outline for a functional taxonomy of document annotation. The question that structures the construction of this taxonomy is “what is the annotator doing?” — that is, we classify annotations according to their function, objective, or purpose.

Status: This document is an outline for a more comprehensive treatment; we have briefly presented some of the main ideas in order to initiate discussion, but in this document they are not developed fully or explicitly related to existing work.

Contact: For more information about this document or its contents contact:

Dr. Allen Renear, Director
Scholarly Technology Group
Brown University / Box 1841
Providence RI, 02912
v: 401-863-7312 e: Allen_Renear@Brown.Edu

or

Dr. Steve DeRose, Chief Scientist
Scholarly Technology Group
Brown University / Box 1841
Providence RI, 02912
Steven_DeRose@Brown.Edu

The Brown University Scholarly Technology Group (STG) is an applied research and development group specializing in electronic document design and markup systems.

Contents

About this Document	1
Contents	2
Introduction	4
What is an Annotation?.....	4
Nature of this Taxonomy.....	5
Other Dimensions of Classification.....	5
Purpose of this Taxonomy.....	6
Limitations of this Taxonomy.....	7
Taxonomy	8
A. Recording and Scheduling Reading	8
1 Marking Segment(s) as Read.....	8
2 Marking Segment(s) as Unread.....	8
3 Scheduling Segments to be Read.....	9
4 Trails.....	9
Implementation implications.....	9
B. Basic Highlighting	10
1 Non-Specific Emphasis.....	10
2 Specific Emphasis.....	10
3 Non-Substantive.....	11
Implementation implications.....	12
C. Commentary	12
1 Analysis.....	13
2 Paraphrase.....	13
3 Clarification.....	13
4 Systematic Note-taking.....	13
Implementation implications.....	14
D. Classification	15
1 Citation Source.....	15
2 Subject Classification.....	15
Implementation implications.....	15
E. Copyediting/Editing/Joint Authoring	15
1 Copyediting.....	16
2 Substantive Editing.....	16
3 Collaborative Authoring.....	16
4 General Editing.....	16
5 Updates and Revision.....	16
Implementation implications.....	16
F. Speech Acts	17
1 Declaratives.....	17
2 Verdictives.....	17
3 Directives.....	17
4 Acknowledgments.....	18
5 Commissives.....	18
6 Assertives.....	18
Implementation implications.....	18
An Example of Using this Taxonomy	18
Using Annotations	21
Private Annotation.....	21
Shared Annotations.....	22
Implementation Implications.....	22

Appendix: Requirements and Functionality Consequences	24
General Form of an Annotation	24
Special Types of Associated Metadata	25
Author, Time of Creation, and Type	25
Keyword or type assignment.....	25
Attachment point.....	25
Privacy and Permissions.....	25
Due Date.....	26
Signatures.....	26
Annotation features	26
Sorting.....	26
Filtered access.....	26
Searching	26
Deletability.....	26
Import/export.....	27
Action triggering	27
Annotation display	27
References	29

Introduction

All of western civilization is just a series of footnotes to Plato

— Alfred North Whitehead, on the evolution of western civilization

... I would always grab the dirtiest copy of a book from the library ... the dirtiest ones had the most marginalia...

— Andries van Dam, on the evolution of annotations in FRESS (Hypertext '87 Keynote Address).

Annotation is a fundamental and pervasive feature of knowledge-work and communication, one which plays important roles in the workplace, the classroom, and the broader world of culture, commerce, and entertainment. The importance of annotation is evident enough in non-digital activities, but it takes little imagination to see the vast opportunities, both for new applications and for new functionality, in computer-based tools for communication and information management. It is not surprising that almost every experimental computer-based system for supporting communication and knowledge-work that has ever been imagined or built has had high-function annotation as a central feature (Bieber and Vitali 1997). But digital documents have been both bad news and good news for annotation practices. Good news because the opportunities for enhancement seem so extraordinary, but bad news because current technology has not fully lived up to this promise. We hope our work will throw some light on why this is so and suggest some new strategies for future directions.

What is an Annotation?

The term *annotation* can be applied to many activities and text features in traditional media. Readers need to perform similar activities in digital media. As many researchers who have investigated annotation in digital media point out, this integral part of the reading activity has to be supported in a manner that is as unobtrusive and free-form as it often is on paper. In this paper we focus primarily on readers' interventions that are anchored to some point or span in the text or other information resource, and are typically linked to brief markings or commentary. There are some other important types of annotation in traditional media, such as "out-of-line" notes, in a notebook, with only text fragments or page reference to indicate what they are an annotation of, that we will try to take into account.

To begin with, digital annotations are created in similar fashion and for similar reasons as traditional annotations, except that the electronic medium provides added functionality to the reader and author. But once the initial hurdle of providing digital annotations in a usable way has been passed, annotations in ebooks or on other digital documents can be manipulated, searched, ordered, filtered, and shared, without the reader having to flip through the pages of the book or having to try and correlate information on a piece of paper via an alphanumeric readout on a VCR or tape deck.

Digital annotation can instantiate and extend traditional mechanisms like the footnote using link mechanisms. In this case, the annotation may consist only of the link itself, although typing information and commentary can also be added. Digital annotation can also be added to documents in media other than text, both still (images) and time based, and is itself not restricted to text, but can be audio, video, or graphic.

Although some activities, such as trails or paths, simple bookmarks, or threaded discussions, may not seem to belong in the category of annotation, we have erred

on the side of inclusiveness wherever we thought we could identify significant common features and discuss how these activities fit in.

In this document, we are discussing annotations that can be used on many sorts of hardware, software, and digital documents, and not limiting ourselves to text annotations on textual documents. For example, the categories and functionalities apply equally to eBooks, handhelds, desktops, and distributed clusters of documents like the Web, or an intranet.

Nature of this Taxonomy

In order to talk about what annotation features should be supported in what MS products, we felt that it would be best to proceed top-down, starting with those annotation activities (purposes, functions) that could be identified as essential. Once a (growing) set of requirements is agreed upon, we can then map annotation activities to mechanisms that support them, and then map the set of mechanisms to various hardware/software platforms. For example, active readers use (colored) highlighters as a mechanism to show *emphasis* or *importance*, *definitions of terms*, *agreement or disagreement*, *progress through a document*, etc. Post-its are used to flag *to do's*, *places to (re)turn to*, *record commentary* without marking up the original, etc. Thus multiple, functionally separable annotation activities or purposes can be accomplished via the same mechanism, and therefore the mapping from annotation activity to mechanism is many to one, or even many to many. Similarly, a mechanism may have different implementations in different applications on platforms with different form factors, again a many-to-many mapping. Clearly there is a tension between wanting to have an “economy of mechanism” (to minimize the learning and implementation efforts) and lack of ambiguity in having a unique mechanism (with a unique representation in the document) for each important, distinguishable kind of annotation.

However, annotation facilities can't be designed strictly top-down: ultimately implementation details affect functionality. For example, if all that implementers have available at the platform level for anchoring annotations is HTML URLs, then the functionality of the annotation mechanisms that can be implemented will be severely limited (this is one reason annotation has not yet become widespread on the Web). Conversely, the XML-based, XPointer/XLink protocols will enable much richer and more useful annotation mechanisms and functionality.

This paper starts the top-down process by presenting an outline for a functional taxonomy of document annotation. The question that structures the construction of the taxonomy is “What is the object or purpose of the annotation?”

Although there have been various attempts to classify annotation and annotation systems based on various features and behaviors, there have been few efforts to develop a classification based on the *purposes* of annotation. (One of the important exceptions is Marshall, for example 1997; we make use of her work here.) We think that such a taxonomy is critical not only for understanding annotation in general, but, as described below, for analyzing the technological issues and market opportunities, and for designing annotation systems. (For an important related project, a typology of reading goals carried out at Rank Xerox Research Center, see O'Hara 1996.)

Other Dimensions of Classification

There are many dimensions other than function along which annotations can be classified (Marshall 1997, 1998a, b, Phelps and Wilensky 1997, and others). We note four basic dimensions which we believe are particularly important and which complement our functional classification:

- 1) *Markets or domains of practice*. Annotations may be classified according to the different market domains in which they are used, such as consumer reading, technical document delivery, publication support, workplace

collaboration, education support systems, and web forums and communities

- 2) *Mechanisms and Features*. Annotation mechanisms may be organized into categories (revision bars, stickies, highlighting), along with their specific features (multiple or single attachment, typing, access permissions, display options, filtering support, structured text support, shareability, etc.).
- 3) *Media or Technology Platform*. Annotation systems may be classified according to the sorts of media and data that are being annotated. For instance, desktop publishing and word processing documents, web pages, still images, data and model driven interactive illustrations, streaming digital video, streaming digital audio, etc.
- 4) *Products*. Annotation systems may be classified according to the particular software and hardware products that support annotation: eBook, MS Internet Explorer, MS Office, MS Windows, MSN, etc.

Truly useful insights into design issues for annotation support require all these areas to be illuminated, as we wish to ask questions that cut across them — such as “what permission features would be found valuable by teachers using IE to annotate streaming audio?” Although our main task here is categorizing annotation according to function, we frequently make comments on the relationships between our functional categories and markets, features, media platforms, and products.

Purpose of this Taxonomy

The motivation for developing this taxonomy is to facilitate the design and development of software systems supporting annotation. Its value will be determined by its usefulness in supporting several sorts of analyses:

1. *General Market Requirements*. What sorts of annotation, understood in broad functional terms, should be supported in different market domains? For instance, do ordinary trade book readers require support for “Basic Highlighting”? (yes, a nearly absolute requirement); do they require support for “Copyediting”? (no, of no interest whatsoever); do business systems require support for “Speech Acts”? (not a requirement, but might be valued if done well); and so on.
2. *Specific Market Requirements*. This analysis would pose more specific questions about exactly what *features* are required when certain sorts of annotation are supported. For instance, do annotation systems supporting Commentary annotation require date / time stamps, and annotator fields, in workplace applications? (of course they do); do they require those features for use in general consumer reading? (not necessary, but of some value). Must Commentary annotations be shareable in the workplace? (absolutely); in general consumer reading? (no). Do educational applications of annotating video require filtering? (yes), sharing? (would be of considerable value); does consumer conferencing require signature mechanisms for Speech Acts? (probably not, unless for commerce). The question here is what opportunities for new valued functionality are presented by the combination of software capability and user goals and expectations in different market domains.
3. *Software Functionality Requirements*. What sorts of annotation functionality should be supported in different software products: eBook, IE, Office, MSN Web Communities, etc. given the capabilities of these products and their market position and strategy? I.e., what do current or intended users of these products expect and value, what functionality *would* they value?

4. *Mapping to Basic Mechanisms.* What are the basic mechanisms that will support the various annotation activities identified in our taxonomy? How can these basic abstractions be implemented in core technology to permit code sharing if/when possible?
5. *Interface Unification.* Given a basic set of mechanisms, how can we create a common appearance and set of behavior expectations across different software products and their user interfaces?

A paper developing this taxonomy further, and connecting it explicitly with the business and software development problems described above, will be prepared. In the present document, as a very simple exercise, we show how our taxonomy can be used to organize a General Market Requirements matrix in order to assist in asking and answering the questions in that category; similar applications of the taxonomy to solving problems in the other categories will be presented in the final paper.

Limitations of this Taxonomy

Given the goal of developing a taxonomy that provides practical illumination of issues relevant to developing business strategies and software systems, and our interest in developing a useful practical taxonomy in a reasonable amount of time, we have accepted two specific limitations:

- 1) *Categories reflect actual clustering of practices*

A pure theoretical taxonomy would always consistently deploy rigorous logical principles of classification, even when doing so diminished the usefulness of the results (by increasing their complexity and obscurity). We instead allow departures from such principles when doing so makes for a simpler and more intuitive scheme. We note that efforts in other domains to begin with theoretically rigorous analyses (e.g., thorough-going applications of speech act analysis to work-flow and conferencing) as often as not fail to produce useful systems or practical insights in a reasonable amount of time. Purists who are offended by this compromise might consider the process iterative: future work may continue the process of analysis and distinguish the dimensions that we currently conflate (following actual annotation practice and user models, rather than classification principles alone).

- 2) *Cases and categories are mainly drawn from actual practice with, time-invariant documents*

In this first document, we draw most of our prototypical examples, and therefore our categories, from actual annotation practices with physical documents. Although we consider all kinds of physical documents, from paperback novels to technical manuals, to business documents, to the scraps of miscellaneous printed and written material that clutter our lives, we have not yet fully considered time-based media or digital media — even though we have those applications in mind (Barger et al. 1998, Goerwitz 1996, Sawhney 1996). We believe that this approach actually makes for a clearer and more compelling initial taxonomy, and one that is, in fact, nearly complete in any case — the basic sorts of things humans want to do when they annotate turn out to be relatively stable and invariant across media and technology platforms. But we do know that new media (e.g., streaming video), new functionality (e.g., collaborative filtering, sharing, audio annotation), and new market domains (e.g., consumer computer conferencing, consumer video entertainment) will almost certainly yield at least some new functions that we will need to accommodate.

Taxonomy

Our classification of reasons *why* people annotate begins with six main categories of annotation. Each of these areas encompasses several more detailed types of annotation, but within each category, there are similarities of functionality and use, that lead to commonalities in how they may be implemented and manifested to the user via the user interface. The major areas we distinguish are:

- A. Recording and Scheduling Reading
- B. Basic Highlighting
- C. Commentary
- D. Classification
- E. Copyediting / Editing / Joint-Authoring
- F. "Speech Acts"

A. Recording and Scheduling Reading

One of the most common uses of annotation is to indicate that some text has been read, or to schedule text for reading (or for not reading). Its specific applications range from simple bookmarking to scheduling assignments, to designing complex reading paths, to creating abridgments; in addition there are some odd and neglected (although common) practices that must be supported if the digital reading experience is to be truly "bookish". We divide the application of Recording and Scheduling Reading into four more specific annotation purposes:

1. Marking Segment(s) as Read
2. Marking Segment(s) as Unread
3. Scheduling Segments to be Read or Omitted
4. Trails

1 Marking Segment(s) as Read

Here the reader is simply indicating that text or other material has been read or viewed. In print books a tick mark, bent page, or a foreign object — the traditional "bookmark" is common. A more powerful but similar marker is the commonplace vertical line indicating a section read or horizontal underline indicating lines of print read; unlike bookmark-like devices these do not assume continuous linear reading.

Some readers (particularly students, but also professionals reading technical materials) sometimes trace their whole reading history with a highlighter or by underlining. Sometimes such marks seem to be little more than doodling, rather than a placemaker, and it is hard to otherwise explain this surprisingly common practice except as at least in part an aid in concentration, and a way of ensuring active engagement with the content, rather than being a genuine act of leaving information in the book.

2 Marking Segment(s) as Unread

Here the reader wishes to indicate material that has *not* been read. The familiar bookmark described above does this as well, as it partitions the book into read and unread portions. But in addition the specific indication of segments as unread is also common. In print books this may be done with paired marginal marks, a vertical line along a passage of text, etc. The words "not read" or "unread", or special idiosyncratic symbols may be used. Unlike traditional book marks, such techniques can easily mark non-contiguous segments of unread text.

The practice of marking text as *to be read* (discussed below as “scheduling”) in a sense implies marking that it is *unread*, although we must accommodate the (common) case where the text has been read but is being scheduled to be read again.

3 Scheduling Segments to be Read

Here the reader (or annotator) is marking material to be read, or to be left unread, or scheduling what will be read when. Often the annotator intersperses reading assignments or goals referencing other documents, or even other non-reading activities (“practice the double banked shot before reading the next section”, “see class notes on differentiation first”, “read documentation of how this is put together before continuing”, “go see TA”, “do experiment”, “assign this to Smith”, etc.). In these instances, we are not scheduling reading, but integrating a reading schedule with other closely related activities. (We also note that sometimes readers indicate activities entirely irrelevant to the reading — e.g., “eat lunch”, “set alarm”, “do ten pushups” etc.; sometimes as a reward, sometimes as a reminder.)

4 Trails

Although trails, or guided reading paths, relate to several of the taxonomic categories, they fit particularly well here, since they serve to lead the reader through a sequence of locations in a document (or more likely, several documents). At their simplest, trails are merely lists of “to read” annotations; generally with a name, and a mechanism for treating the list as a unit distinct from other lists or other individual annotations. However, trails are vastly more useful if they also support at least brief commentary: why this section is relevant, how it is related to neighboring ones, cross-references off the path, and so on. Such trails are standard equipment in collaborative work systems, but also in technical manuals, policy and procedure books, instructional and research material, and other genres.

One of the main differences between traditional annotation and paths or trails is that the former is strongly attached to the text it is annotating, and tends to be a small addition to a large pre-existing resource or document. Paths and trails are substantial new documents themselves, and often incorporate small parts of the original resource or document, leaving the rest aside (Zellweger 1989, Mylonas and Heath 1990). Trails were an important part of Vannevar Bush’s vision of hypertext (Bush 1945). Probably the earliest electronic community to make extensive use of filtered trails was the the NEH-sponsored poetry experiment conducted at Brown in the mid seventies using Fress (Catano, 1979).

Implementation implications

Recording and scheduling readings requires very little of the annotation system. A user interface for tracing reading may be very valuable to the reading experience, even though such tracing is of little general use later, and should be easily deleted so it does not clutter the limited display space. Scheduling things as “to be read” merely requires a general way to indicate a range of content (often, but not necessarily, units such as paragraphs or sections), a way to recover such annotations in order to actually *do* the reading, and ways to delete them or mark them “done”. Using a surrogate like the table of contents is also common; if the units to be read or ignored are predefined like chapters or sections, marks could be inserted there as well. Since there are several types of marking, distinct icons or catch-words, and a way to step through annotations by types, quickly become critical. Assigning due dates in the manner typical on pocket organizers is also useful, especially for school and business users. For trails, of course, a way to step forward and backward along the trail, and a simple way to export and share trails with others, are critical (these very basic functions are, of course, also applicable to many other types of annotations).

B. Basic Highlighting

Here the reader marks a bit of text for some reason or other. Underlining passages is common, as are vertical lines (often repeated for intensification), and other marks. At schools and universities translucent ink marking with highlighters is common. Sometimes the reason for marking a segment of text is indicated cryptically and idiosyncratically with a special symbol, and sometimes with common and more publicly meaningful linguistic or para-linguistic conventions, such as punctuation, phrases, or words. When highlighters are used color-coding is not uncommon.

Obviously at one extreme it is hard to distinguish basic highlighting from simply recording that a passage was read; and from another extreme (extensive comments) it is hard to distinguish Basic Highlighting from Commentary or Classification. But for now we think it is best to break Basic Highlighting out in this way: Basic Highlighting is a relatively common practice where readers mark certain passages to express a few very basic attitudes towards them, often with non-linguistic or para-linguistic marks, sometimes with a word or short phrase. Commentary and Classification differ from Basic Highlighting mainly in degree of complexity, and are far less common in general. Commentary involves more lengthy and structured linguistic annotation; Classification usually implements a more systematic analysis and, typically, the use of standardized categories (as in indexing, whether professional or the ad hoc lists on the fly leaf of relevant page numbers where a topic is treated). Public or shared systems of classification are important for cataloguing documents (as librarians do), but are rarely if ever used *within* a document to annotate textual passages.

We divide Basic Highlighting into three varieties, distinguished by the nature and specificity of the annotator's objective:

- 1) Non-specific Emphasis,
- 2) Specific Emphasis,
- 3) Non-Substantive.

1 Non-Specific Emphasis

Here the text is marked to direct attention to it, but the annotator has no further notion, or no clear notion, of exactly *why* it is being highlighted. Such an annotation cannot always be clearly distinguished in intent from annotations marking a segment as read (above), or, on the other hand, segments marked for a reason, as below. Or, the distinction may be clear but only to the annotator. Non-specific emphasis is typically much too vague in intent and idiosyncratic in execution to be intended to convey anything to anyone other than the annotator; it serves more to remind him of his state of mind when he first read the passage and to record that he has indeed read the passages. This sort of annotation, Non-Specific Emphasis, is *extremely* widespread through all sorts of reading and must certainly be supported by any system which attempts to maximize similarities with ordinary reading.

2 Specific Emphasis

In cases of Basic Highlighting with *Specific Emphasis* the annotator is highlighting a passage in order to record a specific attitude or appraisal. Annotations of specific emphasis are still typically intended primarily for the use of the original annotator upon a later occasion, but in many cases the meaning is clear enough to others, and the annotator may well intend to share his approval, confusion, etc. This sort of annotation is extremely widespread and also must be supported by any system which attempts to maximize similarities with ordinary reading. Also, readers are very used to inventing their own simple symbol systems, so locking down a fixed set of emphasis types will detract from their experience. In some cases (bordering upon Commentary and Classification) some readers

engage in practices such as collecting page numbers of such marks in the front of their book — opportunities for supporting this functionality in digital devices is obvious.

Important

The text is marked as important. This may be simple, or qualified as to the sort of importance, relevance, etc. This is almost certainly the single most common sort of annotation in widespread use. Underlining, repeated vertical lines (the number of lines representing the degree of importance), corners, translucent ink highlighting, and exclamation marks (again, with repetition indicating degrees) are all in common use for marking text as important.

Agreement

The annotator agrees with the author. Extremely common. Check marks, exclamation marks, words like “yes!”, are all used, with degrees indicated by boldness of line, size, repetition, and so on.

Disagreement

The annotator disagrees with the author. Extremely common. “No!”, “?” or “Q” (for “I question this”) are common marks, again often modified to indicate strength of feeling. (Of course specific vehement and scatological comments on the intellectual quality of the writing or intelligence or motivation of the author are not uncommon accompaniments.)

Confusion

The annotator doesn’t understand the passage. Extremely common. A question mark, sometimes repeated, sometimes with suffixed or interspersed exclamation marks are common signals for confusion, as is “Q” (for “I have questions”) or the phrase “huh?”, etc.

[other]

In addition we note that annotators frequently have other very specific clear reasons in mind. For instance, one fairly common intention is to mark something as “quotable”, either for public presentations or future writing.

3 Non-Substantive

Here we take up some sorts of Basic Highlighting that are specific enough for the most part, but are distinctive in not being responses to the *content* of the text or the communicative intentions of the author.

Linguistic Uncertainty or Strangeness

Here the annotator is confused about the meaning of a word, the antecedent of a pronoun, the syntax of a sentence. A variety of marks are used to indicate this uncertainty and to represent alternative resolutions, including links implemented by circles and arrows, page references in the margin, and so on. This is fairly common in all annotating; many readers who are trying to improve themselves by reading would be extremely frustrated if they could not circle unusual words, either in order to look them up later or to help them remember them so they can try them out in conversation. And, more positively, they might be pleased to classify the annotation as a “new word”, so that they could collectively review discovered new words at some later time.

Incidental Connections

Here the passage is of interest to the annotator, but for reasons having nothing to do with the intent of an author, or the interest of the reader in understanding or making use of the material. For instance, someone reading technical documentation in order to install software marks a sentence because its syntax is remarkably elegant — or remarkably ugly. Or someone reading a novel might mark a character's name because it happens to also be the name of the reader's relative. Although one rarely thinks of these as annotation practices, they are very common.

Another case of annotations marking incidental connections is when the reader records in the document aspects of the occasion when he was reading the document or passage. For instance, the exotic locale, the special friend, the meal, the weather, etc. – and, of course, almost always the date.

And a final case occurs in books given as gifts: The giver often inscribes a blank page near the front, with a personal note to the recipient. (In the case of eBooks, which may be given as gifts, this capability of printed books would be a shame to lose, and may present opportunities for enhancement.)

Implementation implications

The basic facilities for highlighting are largely the same as needed for marking and scheduling reading: a way to indicate scope, and to attach type-of-annotation signals. The types become even more important, however, and users certainly need to be able to coin their own. One simple way to do this is like item categories in pocket organizers and file systems: users can create a small set of named categories, perhaps each with a corresponding icon, color, or other sign, and then easily attach the categories to selections. A small set of initial categories can be helpful to get things started. So to create such an annotation, perhaps the user need only select and then pick from a menu of types. The type can be indicated by an icon, and simple facilities can allow stepping through annotations by type. This buys quite a lot of power and user freedom with a fairly simple mechanism.

Idiosyncratic inklike annotation is particularly important to booklike reading. It can take the form of circling words, selecting irregular sections of text and making graphical annotations in the margins. Its value is not only aesthetic, although users find that important; the actual geometry serves as a memory aid and free-form drawing capabilities allow the user to develop personal categorization schemes and associations. For this reason, it is important to be able to preserve inklike markings as they were written, even if the document is reformatted. For more on this see Appendix A under Annotation Display.

C. Commentary

When an annotation contains more than a word or short phrase or represents an analytical representation of the content of the document we call it *Commentary*.

Historically and currently, commentary is extremely important but its use is constrained primarily by the limitations of the medium containing the text to be commented on and the difficulty in establishing connections with other media. In the case of a book one has very limited space on the pages for much prose commentary (end pages, fly leaves, etc. are also fairly limited, though commonly used). In the case of reports printed simplex, one of course has the verso pages for Commentary; some important works can be purchased in special editions with wide margins or blank verso pages just to facilitate such annotation.

Extended commentary is perhaps not considered a core aspect of the popular book reading experience for most people (students and scholars being the

obvious exceptions) but it is more common in reports, drafts, manuals, etc., than one might think at first.

Much of the annotation found in network conferencing, whether business or consumer oriented, is probably best considered a form of Commentary, given its tendency to evolve into independent textual passages as well developed as the text being commented on. Moreover in general conferencing and threaded discussions there is no specific text or texts (other than the preceding contributions) that is being “annotated” in the narrow sense at all, although indefinitely many other documents, and other information resources (network accessible video, sound, etc.) may be referred to with annotation-like links. Threaded discussions may be considered one point on a multi-dimensional spectrum of annotation and the different practices and mechanisms that are particularly important there will differ from the ones of most importance in single-document centered annotation.

Note: Commentary-style annotations obviously present some of the most important opportunities for supporting reading and writing. For instance, the ability to manage annotations intelligently in a separate annotation space, and to filter and share them, are extremely compelling ideas, dating back at least to FRESS (Carmody et al. 1969, DeRose and van Dam 1999). A particularly innovative recent idea is implemented in Xlibris, which automatically tracks user annotations and creates and executes queries based on an analysis of their content in order to maintain a visible list of relevant passages from other documents (Schilit 1999).

Commentary can be divided into several types:

1 Analysis

A prose discussion of the passage, perhaps expressing agreement, disagreement, confusion, application, connection with other topics, etc. In print books it often accompanies a mark of Specific Emphasis.

Note: This kind of annotation is one where two-ended attachments are important. If you’re analyzing (whether for academic reasons, for picking stocks, for keeping notes as you study a policy manual, or whatever), you will likely think of relationships to other things you’ve read, and want to find them or link them in somehow. A simple way to create such a connection is invaluable.

2 Paraphrase

Paraphrase, summary, or abridgment of passage, particularly important in professional reading and educational settings.

3 Clarification

Related to paraphrase, but with the specific purpose of making the text more understandable (to the same reader, or another). Arguably translation is an instance of clarification. Another practice which is difficult to locate in our taxonomy but which we’ll mention here, consists of various kinds of calculations and “worked problems” (Marshall 1997).

4 Systematic Note-taking

Strictly speaking this is not a separate category, but rather a special practice which is a particularly intense application of the other categories. Like lengthy commentary above, it pushes the conceptual bounds of the very notion of annotation. Here the annotator follows a more or less highly developed system of note-taking, creating an organized set of notes, classified as to topics,

documents, authors, etc., and often containing a mixture of paraphrase, explanation, classification, quotation, etc. Over the years many different and complex systems for note-taking have been developed, and several related systems (typically taught to high school students and using 3x5 and 4x6 index cards) have been very popular this century. For many years now there have also been database systems specifically designed to support such note-taking, although users are more apt to use their own (usually loose-knit and haphazard, sometimes coordinated) systems composed of word processing documents, operating system directories, email notebooks, URL bookmarks, bibliographic databases, etc.

We note that as attractive as the prospect of computer-supported Systematic Note-Taking is, it has proven to be a tar-pit of failed initiatives and programs, both analog and digital, at least as far as commercial and educational applications are concerned. Many have seen a market here, few have found one. Far more likely to be successful is starting at the low end and scaling up the functionality of basic annotation and communication systems.

Implementation implications

Commentary annotations of all kinds share the characteristic that they attach actual text, not just a terse label such as one might pick from a menu. This requires that some means of text entry must be possible in the short term; in the long term, commentary should support other media as well (some platforms allow simple graphical and audio annotation already). The issues of annotation type are much the same as for prior categories.

Commentary that effectively proposes changes to the text, such as paraphrase and clarification, may benefit from being displayed directly inline, while not obscuring the “official” or original text — say, as square-bracketed phrases (a familiar convention for text insertions), or block-quotations with some particular formatting. If the paraphrase or clarification is a good one (and presumably its creator often thinks it is), seeing it automatically the next time around may be a very pleasant option — as well as mimicking the paper reading experience better than an on-demand dialog box. Even if commentary has to be pulled up on demand rather than appearing automatically, having it appear inline avoids the sticky problem of whether its window or dialog box obscures the actual text being commented upon (Zellweger et al. 1998).

Commentary should be kept in XML markup. But there is no need for it to conform to the document's DTD (and probably no advantage). Certainly, it is not parsed as if it were an integral part of the document, and need not be valid in the location where it is attached. A very simple DTD just for annotations is probably enough for 90% of uses. Anyone doing extended commentary will need to use emphasis, quotations, lists, and such, and probably also editorial elements (sections, subsections, headings, etc.). Domain-specific constructions (name, price, part number, equations, citations, etc.) may be needed at least in technical and commercial writing almost immediately; but these come closer to the category of editing (E), discussed below.

Since commentary annotations consist of extended text, they provide a good example of annotations that can be made “in-line” in traditional documents, for example in the margins, or “out-of-line”, in a notebook, with page numbers or other types of link indicator. In digital media, it is important to be able to display this kind of annotation in both ways, either as discrete comments which come up at a particular place, or as a document consisting of comments. These comments would each be linked to their place of origin in other document(s).

D. Classification

The above categories might all be very broadly considered “classification” in some way or other, but actual reader practices seems to make it more appropriate to have a separate category for annotations whose fundamental purpose is to classify the passage according to some scheme or other. We distinguish two sorts of classification:

1 Citation Source

Here there is a specific identification of the source of some quoted or paraphrased text or idea. This might be done either because it is unstated in the passage, or to give a standard form or additional information for later use. For example, pithy quotes may only cite the author without giving the source text, or may be used with no attribution at all (which is why one often sees inquiries posted on various email lists, asking for the source of a given quote seen somewhere); indicating the source, or full source, is a common annotation practice, and one with interesting and well-known potential for inter-document hypertext links.

Note: A way to post such corrections and amplifications to a publisher’s Web site could be of interest to readers, and would help publishers capture and distribute value.

2 Subject Classification

Here the annotation is a bit more analytic, and renders an opinion about parts of the text and how they fit into a larger organization scheme. Some schemes are minimal and informal, others are elaborate and rigorously developed. Often personal idiosyncratic classification schemes are used. General book-indexing is itself an instance of this sort of annotation — sometimes using shared domain-specific subject descriptors, although these are rarely seen below the level of monograph or article classification. Obviously software support for subject descriptors would present many well-known opportunities for valuable new functionality, particularly in the areas of retrieval, navigation, and information management.

Implementation implications

Classification is a more structured kind of annotation, probably requiring forms for input and more detailed XML markup for the underlying data representation to be most effective. In the short term, the implementation can be reduced to that of the previous category, though the added value to the user is less: for example, users cannot export their classifications or references so easily for re-use by other software if these are not represented in a structured way. Stronger support for classification could relate to sales and marketing strategies for publishers: making cross-references easy and machine-processable could provide added sales opportunities, as the referenced objects become available online.

E. Copyediting/Editing/Joint Authoring

Although copyediting, correcting, revising, editing, and joint authoring are not part of the ordinary book-ish reading experience, they are one of the most common and important sorts of annotations one makes in the workplace, where one is frequently a participant in document development.

We distinguish five sorts of annotation activities along these lines:

- 1) Copyediting,
- 2) Substantive Editing,

- 3) Collaborative Authoring,
- 4) Updates and Revision,
- 5) General Editing

1 Copyediting

This is the familiar correction of spelling, punctuation, grammar, missing words, formatting and layout, etc., and may include queries to the author on these topics. It is often done on paper using fairly standardized “proofreaders marks”.

Note: One key distinction is that between changes to the text itself, and comments or “side notes” to the one who would do the corrections. Systems supporting annotation must help with this.

2 Substantive Editing

Proposals for addition or deletion of phrases, sentences, passages, citations, etc. These may include substantive queries to the author on analysis, rhetoric, etc., and small to large suggested replacement text or other media.

3 Collaborative Authoring

Proposals for addition or deletion of any material whatsoever. May include substantive queries to co-authors on analysis, rhetoric, workflow, negotiation of responsibilities, schedule, etc.

4 General Editing

When once one or more reviewers have created various editorial annotations, many publishing workflows send them to a general editor who must combine and execute them and then confirm that the results are as intended.

Note: That person will probably be getting contributions from multiple sources, and will need not only to combine them, but to annotate them as to status, priority, author, processing, etc. For instance, as they work they may want to flag various annotations as “taken care of”. And at the end, they must be able to delete all the annotations at will, without manually navigating to each one and executing several commands. Opportunities for software tools to support this process are obvious.

5 Updates and Revision

We include here the case where an annotation-like process is used to update or revise a document, either by a technique that makes a change, obliterating the original material; or one that adds supplementary material that corrects or obsoletes the original.

Note: This is as good a place as any to note that annotations must gracefully relocate themselves as the underlying text is being revised. At least a warning must be issued if a version change is noted. Better would be to use the XML linking checksum facility (for a description, see <http://www.w3.org/XML/Group/1998/12/NOTE-xptr-rqmts> in the members only portion of the w3c website) whenever a version change is noted to see if the document portion annotated was changed or not.

Implementation implications

Editing, like classification, is a more structured kind of annotation. One way in which it differs technically is that the annotations are likely to include recommended changes, which are most naturally expressed in the same tag-set as the document being annotated (classification annotations tend, instead, to be

structured in consequence of their own type, such as bibliographical references). In an ideal fully integrated review application, reviewers could provide suggested corrections that could be validated against the source DTD; they could see their changes in effect, temporarily replacing the original text, or by toggling back and forth. Such an environment also requires smooth integration with authoring systems, with management facilities for receiving and merging multiple sets of annotations, accepting and rejecting them, and tracking their disposition. Much of this is beyond the short- to medium-term scope of ebook annotations but is clearly related to uses of annotation in other MS product areas, and so should be considered at least to the extent of facilitating interoperability and data sharing.

F. Speech Acts

Linguists use the term “Speech Acts” to refer to the various sorts of acts that can be accomplished by a linguistic expression. For instance, by expressing certain sentences in certain contexts you can accomplish such acts as agreeing, disagreeing, approving, hiring, firing, warning, promising, marrying, commanding, making an offer, accepting an offer, and so on. Annotations expressing such acts are extremely important in the workplace — arguably the sole purpose for developing and circulating documents at all (excepting technical documentation) in business and commerce is to elicit acts of this kind (approvals, agreements, etc.) from the appropriate people.

The theory of Speech Acts is a powerful one that could be used as the fundamental principle for organizing a functional taxonomy for annotations. Unfortunately, the categorization of speech acts is a complicated business with a long history in computer systems design but limited definite results, either theoretical or practical. However we do find it useful to identify a major category for annotations where the so-called illocutionary force of the annotation is its most salient feature. Many speech act taxonomies have been proposed; our categories are taken from Bach and Harnish (1979), although we use the more familiar names developed by Searle (1979) and introduced into the information systems design by Winograd and Flores (1986). For a recent application to workplace systems see Auramäki (1988).

1 Declaratives

Declaratives are annotations that establish new institutional facts with respect to the document or its contents. For instance, they approve, appoint, dismiss, command, fine, levy, vote for, nominate, resign, recommend, enact, annul, veto, etc. Typically the annotator must hold an authoritative role or status with respect to the particular action being made. These annotations are very common and extremely important in the workplace.

2 Verdictives

Here there is an official finding with respect to a document or its contents: “the information is correct”, “this contract fails to meet the indemnification requirement”, “this proposal would violate rule 22”, “the claim is invalid”, “this project will cost \$100,000”, etc.

3 Directives

Directives express an attitude with respect to someone’s doing or not doing some action. Some important examples are:

Commands: Here specific instructions are issued, such as “buy one of these for the shop”, “reply, saying thanks but no thanks”, “organize a group to solve this problem”.

Routing: Here general instructions are given as to where (to what institution or department) the document should go next. Obviously very important to workflow management; less so in popular contexts.

Processing Qualifications: Typically these would be cases where instructions in the text were qualified by further directions given as annotations: “Watch it! pour the dye *slowly*”; “don’t use the carbonized form”, “ignore the clicking sound”, etc. Obviously critically important in technical documentation and process control manuals.

Warnings: This is where a reader includes such useful information like: “Be careful, the mixture will now be flammable!”, “This will discolor rayon!”, “this process never works on Sundays for some reason”, or “When I turned the dial to the left, just like it says, the damn thing exploded and burned a hole in my desk”. These annotations are prevalent, and, obviously, important.

4 Acknowledgments

These are expressions of apology, thanks, congratulations, gift-giving. Very common not only in books given as gifts, but also surprisingly common (particularly apologies and thanks) in office documents and other miscellaneous documents.

5 Commissives

These are annotations which convey promises or offers. Very common as business annotations, in bidding and auction-like environments, and elsewhere.

6 Assertives

Assertives are statements of fact. Although many varieties have been identified, unlike the other Speech Acts they don’t seem much like *acts* at all, and, indeed, most of the examples from our other major categories of annotation would actually be Assertives of one sort or another. However some Assertives do stand out as representing a particular kind of act-like annotation. We mention one of particular importance:

Confirmation of Execution: This is where an annotator indicates that something, such as an enumerated instruction, has been accomplished, and so, for instance, need not be done again. E.g.: “Done, AHR, 9/1/99” or “engines already fueled as described, go on to item 7”. Obviously essential for real-life use of instructions and procedures.

Implementation implications

The distinctive characteristic of speech act annotations is that they have the force of action. Because of this, they impose requirements for authentication and logging that may not be so important with other annotation types. Export and transfer of speech act annotations must be supported, though that is also true of most annotation types. Otherwise, the mechanisms for implementing speech act annotations will resemble those of Highlighting and Commentary.

An Example of Using this Taxonomy

The usefulness of this taxonomy, when it is in final form, will be in helping understand general market requirements and opportunities and guiding the design of software systems.

The following matrix is intended to suggest one such use of this taxonomy. The ranking in each cell indicates just how important a particular sort of annotation is in a particular domain. For instance, copyediting is not important at all in general

bookish reading, but very important in working with documents in the workplace; Speech Acts such as Declaratives are unimportant in bookish reading, but very important in workplace reading, and so on.

The rankings shown below, although plausible, are not actual rankings — no attempt has been made to actually determine the importance of annotations in various domains — but simply an example of how the taxonomy might be used to represent such a ranking.

<i>Annotation Types</i>	<i>Domains ‡</i>	General Bookish Reading	Work place & Business Support	TechDoc reading	Consumer conferencing and interaction
Example Platforms	eBooks	Windows, Office	DynaText	MSN Web Communities.	
A. Recording and Scheduling Reading					
1 Marking Segment(s) as Read	1	1	1	1	1
2 Marking Segment(s) as Unread	1	1	1	1	3
3 Scheduling Segments to be Read	1	1	1	1	2
B. Basic Highlighting					
1 Non-specific Emphasis	1	1	1	1	3
2 Specific Emphasis	1	1	1	1	2
3 NonSubstantive	1	2	2	2	2
C. Commentary					
1 Analysis	3	1	1	1	3
2 Paraphrase	3	2	1	1	3
3 Clarification	3	2	1	1	3
4 Systematic Note-taking	3	2	2	2	4
D. Classification					
1 Citation Source	2	3	2	2	2
2 Subject Classification	3	2	2	2	3
E. Copyediting/Editing/Joint Authoring					
1 Copyediting	4	1	3	3	4
2 Substantive Editing	4	1	3	3	4
3 Collaborative Authoring	4	1	3	3	4
4 General Editing	4	1	3	3	4
5 Updates and Revision	4	1	1	1	4
F. Speech Acts					
1 Declaratives	5	1	3	3	2
2 Verdictives	5	1	3	3	2
3 Directives	5	1	1	1	2
4 Acknowledgments	1	1	3	3	2
5 Commissives	4	1	2	2	2
6 Assertives	4	1	2	2	2

The numbers 1-5 represent a ranking of importance or value to users, with 1 being most important and 5 being of least importance.

Using Annotations

Readers are so used to annotating paper documents, that the way annotation behaves in eBooks and other digital documents will likely be a major determinant of their “feel”. Paper affords such ease and variety of annotation, that it will be a challenge to make annotation sufficiently convenient and painless in electronic books that it will not detract from the reading experience. The reward for the user, however, of becoming accustomed to using the different mechanisms necessary for electronic annotation is that the electronic medium provides possibilities for *added* annotation functionality, which if well utilized will compensate for any other shortcomings.

Whereas in traditional media annotations are difficult to access after they have been created, and require added mechanisms (colored tape flags, hand compiled indices in the front cover, in the case of streaming media, a separate piece of paper with a time written on it) for retrieval and filtering, the key advantages of annotations in the electronic book over paper, are that (a) they can be retrieved in various orders and combinations with far less pain; and (b) they can be shared without having to give away your whole book. Exploiting these strengths is invaluable, so long as it is done without undue complexity. Up until this moment, this discussion has focused on the reader who is creating annotations. We will now turn to the user of the annotations. In the short term, the author of the annotations will also be the user, unless someone else shares the document. As devices become more sophisticated however, annotations will be shared, merged and used by people other than the author, and may even be metamorphosed into new documents.

We can identify several categories of use and sharing of annotations. A single user may make annotations for his/her private use. A reader might also make annotations for public consumption. In this case the reader would probably be authoring her comments more carefully and at greater length, for clarity. Both types of annotations may also be merged with other annotations, which may be on the same document or on different documents. For example, a private set of annotations may be merged with the same reader’s other annotation sets, and she may then want to perform operations on the whole collection. A set of annotations that were meant to be private may also become public and shared. An example of this in traditional media is used books that have been written in (Marshall 1998a, b). A digital document might be loaned in the same fashion, or annotations that began as private notes may be judged “publishable” by their author, and made public. It is important that annotations on digital media be usable in these various ways, and be easily repurposed by their readers.

Private Annotation

For the individual user, creating annotations should be as comfortable, indeed trivial, as possible. Careful attention to minimizing the number of actions required is critical, as is an intuitive way for users to create and maintain their own annotation labels and simple schemes (on par with what they do now with punctuation and other symbols, highlighter colors, tape-flags, and the like) for categorizing their annotations.

As users’ annotation and document collections grow, management will quickly become important. As electronic books and other kinds of documents migrate on and off the user’s physical media, annotations should not impose an added burden. At the simplest, all annotations could live in a single collection, so they are always there when the relevant document is loaded. If memory limitations make this unacceptable, migrating applicable annotations on and off the platform automatically, when a document or document set moves on or off, seems the bare minimum users will expect.

On paper, it isn’t easy to quickly view all annotations on a book or chapter, all annotations of a certain type. Conversely, if annotations are gathered together on

a separate sheet of paper, it is difficult to move from the annotation to the place in the document where it is anchored. Viewing lists of annotations, or stepping through them by type, and searching them by content, type, and simple bibliographic information such as creation date, are basic expectations that users will have. Indeed, users of traditional media have come up with ways of doing just that by rapid page-flipping and indexing.

Single-user annotation should also have some facility for “cleaning up”. Particularly if annotations are kept in a single collection, discarding all annotations for a given book seems an obvious operation.

Finally, there are advanced functions that can only exist in electronic media, starting from searching the text of annotations and dictionary look ups of marked unfamiliar words, to the automatic query generation based on annotation markings implemented in Xlibris (Schilit 1999), consensus analysis based on many readers’ annotations of a single document (Marshall, 1998a, b).

Shared Annotations

Electronic documents provide an opportunity not previously available on paper: sharing your annotations on a document or resource. This is of great potential market value. In teaching situations (whether in school, or in technical support or similar fields), adding value to manuals by annotating them is extremely valuable. Noting errors or confusing passages in the manual, attaching examples, or even citing related special offers, can be valuable information to share with others. In corporate settings it is commonplace to jot annotations on articles before passing them on, ranging from the “isn’t this interesting variety” on up to critical directives for further action.

Being able to share annotations will, we think, add significantly to the medium’s value and attraction in the popular market as well. A simple protocol for exporting annotations as XML documents (in, presumably, a simple DTD designed for the purpose), and passing them around via direct or network connections, will enable people to engage with the text in new ways, and to interact and build greater interest in the publications. This can only add to the market for the platform and books.

Users who are not authors will want most of the same searching, filtering, sorting and aggregating functionality that is described for author/users above.

Features like this have been available in systems since Engelbart’s NLS Augment (Engelbart and English 1968, Engelbart 1984). In Augment, it was possible to select views based on filtering link keywords and other attributes. In Engelbart’s system links were embedded, so the filtering was essentially filtering on content. A similar feature is present in HyperWave, which allows links to be shown or hidden depending on attributes or key words (Maurer 1996).

Implementation Implications

Annotation sharing needs to be carefully designed in a few critical ways. Annotations need to bear simple information such as their creation date, creator, and so on, in particular so that an ill-advised import can be undone, or so unwanted annotations can be easily grouped and deleted. Also, mundane actions such as re-downloading an updated annotation set from somewhere, should not result in repetitions of the same annotations in one’s local file. And the system should always make it easy to distinguish one’s own annotations from those of others. Given a few such basic precautions, sharing annotations, we feel, can be a major advantage of electronic documents.

The whole idea of being able to share and merge annotations, and to collaboratively annotate a document implies that annotation sets exist as documents in their own right. Also, many of the manipulations we have touched on in this section are similar to the link bases developed originally by Intermedia

at Brown, and then used to a fuller extent by Open Hypermedia Systems like Microcosm and DeVise (Hall 1996, Grønbaek et al., 1997). In such systems, different sets of connections between documents can be generated based on link keywords and metadatas that a user selects to filter on.

Appendix: Requirements and Functionality Consequences

In this section we collect a few recent notes, mostly eBook oriented, related to implementation, operational functionality, requirements, etc.

General Form of an Annotation

Given our earlier description of an annotation as “readers’ interventions that are anchored to some point in the text” our first attempt at an abstraction for annotation might be this:

An annotation consists of:

- 1) an attachment point or span,
- 2) a field for its type,
- 3) content (such as text or other symbols or data, including video and audio)

But further reflection reveals that general formats for actual annotations must be a bit more complicated. Some fundamental characteristics to consider:

The type of attachment matters. There are several common ways users indicate what is being annotated. First, they may underline, which is roughly equivalent to drag-selection in GUIs; second, they may make marginal bars, brackets, and the like, implying attachment to all adjacent text; and third, they may circle or box things geometrically. And finally, they may simply put a symbol in the margin, with the exact scope of application left vague or approximate.

There are two conflicting priorities that need to be resolved:

- 1) When users come back, they probably would like to see the exact same indication they made; this is not just an aesthetic desire, but the exact shape and physical placement of an annotation is an important aid to memory. This argues for storing some representation of the actual geometry of the reader’s marks vis-à-vis the physical rendering they were applied to.
- 2) Users may come back to a subtly (or even completely) different layout: they may have turned their reader sideways, or requested large print because they forgot their reading glasses, or have plugged in a supplemental large monitor, etc. In such cases the exact geometry is unlikely to be the same, but one can hardly discard the annotations for that reason. This argues for attaching annotations in terms of logical selections in the text.

The only solution we see is to store both kinds of attachment information. For example, when the reader circles, find all words of text within the circle, and make that range the scope (this is imperfect but probably adequate). The logical attachment can then be applied regardless of changes in layout, hardware, state, and so on; while the geometric details can be reapplied if the reader is in the same layout state. This does require recording a little extra information with the annotation, to say what the state was: perhaps no more than a requested scale or zoom factor, the display size, and the particular reader and version in use. Long-term, we anticipate smart algorithms that would map the geometry, for example creating a similarly-shaped circle that encloses pretty much the same text even in another layout; but none of that sophistication is needed in the first release.

Several attachment points or several annotation fields may be needed, or at least a means of chaining annotations together. This is critical to the ability to create trails, reading lists, and the like.

Linking is necessary, since users often create connections between parts of their documents, and between different documents; this means there is an object out there that encapsulates some number of scopes (simplest case, two), describes them and their relationship somehow, and provides for access between them; typed and bi-directional linking is also highly desirable — hence the importance of the emerging W3C XPointer and XLink standards.

Simple metadata must be possible to associate with the annotations (here we mean metadata broadly: not just cataloging or publication data, but simply meta-information about the annotation). A simple list of attribute/value pairs is a powerful but simple approach that can be smoothly extended to the future.

The content of annotations (that is, text the reader creates) may need to be structured text, at least of a rudimentary kind. Then again, given that (e.g.) an eBook reader will contain support for the eBook DTD, there is not much effort required to permit reader annotations to use any of the same constructs (perhaps they do not gain any performance enhancements from precompilation, but they will rarely be big enough for that to matter).

Special Types of Associated Metadata

Any approach to these issues must be coordinated with various ongoing standards efforts such as the OCLC Dublin Core project. (Wherever possible metadata should be automatically generated (e.g. time and date stamps, author stamps, document id; perhaps synchronization with other features of the users' software environment, schedules and calendars and other events; heuristic use of inherited values and prompts could also be useful)

Author, Time of Creation, and Type

Annotations need at least their own basic author, type, and time/date stamp information.

Keyword or type assignment

Several types of annotations involve the reader labeling the annotated content as being something: interesting, quotable, erroneous, funny; whatever. Easy ways to do this (e.g., with inheritance filling fields automatically, and intuitive pen strokes in the case of, e.g., eBooks) is important as the task is tedious even if the result is valuable.

Attachment point

This is so much a part of the link as to be obvious; however, we note here that annotations should use some reasonable approach, e.g., ISBN anchors might make sense, or other standard identifiers. Within the document, an XPointer using an ID and tree path is the most robust identifier likely to be available in a short enough time frame to use.

Annotations attached to graphics should be able to attach to specific selections, such as a drawn region or a dragged rectangle. Eventually, of course, some graphics will contain true objects rather than just pixels, and attachments to those objects and their states in space and time will become necessary as well.

Privacy and Permissions

Fine-grained management of permission is important for creating collaborating groups and subsetting huge communities of users; as well as privacy and commerce-related security.

Due Date

Some important annotations relate to future actions, and users will need to be able to note, e.g., due dates (and find annotations by date – see below). Even when the user does not specify a due date, it should be possible to flag (perhaps via “type”) that a particular annotation does in fact have “to-do” implications.

Signatures

Some annotations may have legal or administrative force, and may need to be authenticated.

Annotation features

Sorting

Annotations should be sortable by their simple metadata, such as creation date, which reader made them, and type. For such purposes trails should be treated as single entries that subsume the whole trail, since that’s how users think of them (not as a whole lot of annotations just mixed in with the rest).

Some kinds of annotation attachment are very clear to users, and it can be helpful to actually show them in any kind of sorted annotation list. For example, many documents use numbered chapters and sections. In such cases, allowing a way for users to get at annotations by section numbers, and vice versa, will provide a very paper-like feel and ease of use. In other cases, showing the first little bit of the annotated content can spur the reader’s memory and save time when looking for things in an annotation view of some kind.

Filtered access

It is important to be able to step through one’s annotations by type, creation date, or by location in the document, or (in the case of scheduling annotations) by due date, or to step back and forth along a specific trail. It does not take many annotations before having to step through all of them, instead of just those of relevant types, becomes tedious and very frustrating. Although in one sense the reader must do so on paper, it is much less painful because flipping pages is so rapid; and users who find it annoying can easily jot down lists of pages by type on blank leaves of the book, again diminishing the pain. It is a minor convenience to be able to step automatically through a trail.

This feature is also one that can take advantage of a path mechanism if there is one, since ad hoc paths can be created dynamically by sorting and filtering annotations. These may serve as the beginnings of more carefully authored paths, but may also be just temporary holding areas for browsing annotations.

Searching

Readers should be able to find annotations either by the content they attach to (“where is that part I marked as Important, about the market crashing?”), or by the content of the annotation itself (“where’s that note I made to myself that some passage related to my will?”).

Deletability

Annotations should be deletable. Multiple-selection and a delete command in a sorted view of annotations would likely be enough for just about any reader, even heavily-laden reviewers. Deletion by a combination of attributes should also be possible.

Import/export

A great deal of the *added* value that digital documents can bring to the reading experience, is the ability to share annotations. Right now people can only share annotations by giving away the physical book as well; this is inconvenient and impractical. But with digital media, readers can share easily: either through whatever direct ports eBooks come to have (serial, infrared, etc), or through the larger network. The design, however, is critical: it would not do, for example, to import a set of annotations, find them useless or annoying, and then be reduced to tedious, manual, one-by-one deletion of them; or to import an annotation set twice and have duplicates show up right on the book.

Action triggering

Some annotations may be most effective if they can trigger actions later or be automatically updated (Catlin 1989). The most obvious is that publishers may want to include ads for other books or networked documents in the back-matter, as they often do on paper. It would be ideal if readers could activate such ads and then, the next time they plug into a kiosk, network, or other source, they are offered the chance to purchase the book. Likewise, triggering retrieval of free information via the Web would be desirable, for surely there will be free publications for eBooks just as there are on paper, as well as publisher incentives to visit Web sites, buy books by subscription, and many offers yet to be imagined

Annotation display

Various kinds of display are appropriate to different kinds of annotations: the most basic choice is whether the annotation's content is displayed automatically or on-demand (both are highly desirable). Annotations that have no content, of course need not display any.

For on-demand display, or for annotations without content, some marker needs to appear to let the user know there is an annotation in the first place. This may be

- 1) the original mark made, such as a circle;
- 2) font or color changes or underlining;
- 3) an inline marker (like superscript numbers on paper); or
- 4) marginal icons, following book practices from long ago that have recently reappeared. This form is especially important for several categories where readers have a familiar tradition: long brackets to mark off passages, punctuation marks such as "!" and "?", as well as proofreaders' marks, or in- and out-arrows to indicate links and cross-references.

It may be wise to indicate right on the mark, whether or not there is associated text content.

When an annotation's content is actually displayed (whether that happens automatically or because the user requested it in response to seeing a marker, or from an annotation overview), it might appear in several places:

- 1) inline at the point of attachment (perhaps looking like a specially boxed block quotation);
- 2) actually temporarily replacing the original content that was annotated (though probably still flagged somehow to distinguish it from original text, such as by underlining)
- 3) in a side margin or panel(a less likely option in an eBook than on paper);
- 4) in a footnote bar at the bottom of the screen;

- 5) in a separate margin pane; or
- 6) in a temporary window such as a dialog box.
- 7) As a separate document with all the other annotations on that document or document set.

Some annotations, such as tracing, don't need text content at all, and may just be flagged by a revision bar, or by some modification of the annotated text.

References

- Auramäki, E., Lehtinen, E., and Lyytinen, K., "A speech-act-based office modeling approach," *ACM TOIS* 6(2): 126-152.
- Bach, K., and Harnish, R. M., *Linguistic Communication and Speech Acts*, MIT Press, 1979.
- Barger, D., Gupta, A., Sanocki, E., and Grudin, J., "Annotations for Streaming Video on the Web: System Design and Usage Studies," MSR-TR-98-60, 1998.
- Bieber, M., Vitali, F., Ashman, H., Balasubramanian, V., and Oinas-Kukkonen, H., "Fourth generation hypermedia: some missing links for the World Wide Web," *International Journal of Human-Computer Studies* 47, 1997, 31-65.
- Bush, V., "As We May Think," *Atlantic Monthly*, 1945.
- Carmody, S., Gross, W., Nelson, T. H., Rice, D., and van Dam, A., "A Hypertext Editing System for the /360" in *Pertinent Concepts in Computer Graphics*, ed. M. Faiman and J. Nievergelt. University of Illinois Press, 1969.
- Catano, J. "Poetry and Computers: Experimenting with Communal Text," *Computers and the Humanities*, 13, 1979.
- Catlin, T., Bush, P., and Yankelovich, N., "Internote: Extending a Hypermedia Framework to Support Annotative Collaboration," *Hypertext'89 Proceedings*, 1989.
- DeRose, S. J., "Expanding the Notion of Links," *Hypertext'89 Proceedings*, 1989.
- DeRose, S. J., and van Dam, A., "Document structure and markup in the FRESS Hypertext System." *Markup Languages* 1(1), Winter 1999, pp. 7-32. Cambridge: MIT Press.
- Engelbart, D. C., ., "Authorship Provisions in AUGMENT," in Douglas C. Engelbart, *COMPCON '84 Digest: Proceedings of the COMPCON Conference*, San Francisco, CA, 1984, pp. 465-472.
- Engelbart, D. C., and English, W. K., ., "A Research Center for Augmenting Human Intellect," in Douglas C. Engelbart and William K. English, *AFIPS Conference Proceedings of the 1968 Fall Joint Computer Conference*, San Francisco, CA, December 1968, Vol. 33, pp. 395-410.
- Goerwitz, R., "STG's Re-Implementation of Progressive Cevents" Technical Note at: http://www.stg.brown.edu/edu/HTT96/realaudio/stg_cevents/HTTcevents.html, 1996
- Grønbaek, K., Bouvin, N.O., and Sloth, L., "Designing Dexter-Based Hypermedia Services for the World Wide Web," *Proceedings of Hypertext'97*, 1997.
- Hall, W., Davis, H., and Hutchings, G., *Rethinking Hypermedia: The Microcosm Approach*, 1996.
- Laliberte, D., and Braverman, A., "A Protocol for Scalable Group and Public Annotations," NCSA Technical Report, 1997.
- Marshall, C. C. "Annotation: from Paper Books to the Digital Library," *Proceedings of Digital Libraries '97*. New York: ACM Press 1997.
- Marshall, C. C. "The Future of Annotation in a Digital (Paper) World." presented at the 35th Annual SGLIS Clinic: Successes and Failures of Digital Libraries, University of Illinois at Urbana-Champaign, 1998a.
- Marshall, C. C. "Toward an Ecology of Hypertext Annotation," *Proceedings of Hypertext '98*. New York: ACM Press, 1998b.

- Maurer, H. , ed. "HyperWave: The Next Generation Web Solution," Addison Wesley, 1996.
- Mylonas, E., and Heath, S., "Hypertext from the Data Point of View: Paths and Links in the Perseus Project" in *Hypertexts: Concepts, systems, and applications*, A. Rizk, N. Streitz and J. André (eds). Cambridge University Press, 1990.
- O'Hara, K., "Towards a Typology of Reading Goals," Rank Xerox Technical Report EPC-1996-107, 1996.
- Phelps, T. A. and Wilensky, R., "Multivalent Annotations," *Proceedings of the First European Conference on Research and Advanced Technology for Digital Libraries*, 1997.
- Roscheisen, M., Mogensen, M., Winograd, T., "Beyond Browsing: Shared Comments, SOAPS, Trails, and On-line Communities," presented at the Third International WWW Conference, 1995.
- Sawhney, N., Balcom, D., and Smith, I., "Hypercafe: Narrative and Aesthetic Properties of Hypervideo," *Proceedings of the Seventh ACM Conference on Hypertext*, 1996.
- Schilit, B. N., Golovchinsky, G., Price, M. N., "Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations," *Proceedings of ACM CHI 98*, 1998.
- Schilit, B. N., Price, M. N., Golovchinsky, G., Tanaka, K., Marshall, C. C., "As We May Read: The Reading Appliance Revolution", *IEEE Computer* 32(1), January 1999, pp. 65-73
- Searle, J.R., *Expression and Meaning*, Cambridge University Press, Cambridge, 1979.
- Winograd, T. and Flores, F., *Understanding Computers and Cognition: A New Foundation for Design*. Norwood: Ablex, 1986.
- Yankelovich, N., Haan, B. J., Meyrowitz, N. K., and Drucker, S. M., , "Intermedia: The Concept and Construction of a Seamless Information Environment," *IEEE Computer* 21 (1), 1988.
- Yankelovich, N., Norman, M., and van Dam, A., "Reading and Writing the Electronic Book" in *IEEE Computer* 18 (10), October, 1985.
- Zellweger, P. T., "Scripted Documents: A Hypermedia Path Mechanism," *Hypertext'89 Proceedings*, 1989.
- Zellweger, P. T., Chang, B., and Mackinlay, J. D., "Fluid Links for Informed and Incremental Link Transactions," *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia*, 1998. Pp. 50-57.