

Identifying Useful Passages in Documents based on Annotation Patterns

Frank Shipman², Morgan Price¹, Catherine C. Marshall³, Gene Golovchinsky¹

¹FX Palo Alto Laboratory
3400 Hillview Avenue, Building 4, Palo Alto, CA 94304
gene@fxpal.com

²Department of Computer Science, Texas A&M University
College Station, TX 77843-3112
shipman@cs.tamu.edu

³Microsoft Corporation
1 Microsoft Way, Redmond, WA 98052
cathymar@microsoft.com

Abstract. Many readers annotate passages that are important to their work. If we understand the relationship between the types of marks on a passage and the passage's ultimate utility in a task, then we can design e-book software to facilitate access to the most important annotated parts of the documents. To investigate this hypothesis and to guide software design, we have analyzed annotations collected during an earlier study of law students reading printed case law and writing Moot Court briefs. This study has allowed us to characterize the relationship between the students' annotations and the citations they use in their final written briefs. We think of annotations that relate directly to the written brief as *high-value annotations*; these annotations have particular, detectable characteristics. Based on this study we have designed a mark parser that analyzes freeform digital ink to identify such high-value annotations.

1 Introduction

We read for many reasons: we read for fun; we read for general knowledge; and we read with a particular task in mind. When we read for fun, we seldom annotate; but when we read material crucial to a particular task, our annotations may cover its pages. In fact, a diary study of reading found that over a quarter of the writing that occurs during work-related reading is in the form of annotations, marks on top of existing documents [1]. Annotations may perform a variety of functions; for example, they may emphasize a particular point in the document, express a reaction, compare the text to other texts, record a related idea, or generate action items.

We are investigating methods for identifying important passages in an annotated document; as O'Hara and Sellen suggest, these passages may then be used in related activities such as writing or re-reading the document [6]. Just as readers use the table of contents, indices, and document structure to orient themselves during a first pass through the material, their annotations may be used for orientation during re-reading.

In a past study, we examined how law students read and annotated printed source materials to prepare for a Moot Court competition, a venue in which the students argue hypothetical cases. The previous study used the results of interviews and observations to characterize the students' general reading and annotation practices [5]. At the conclusion of that study, we collected the law students' written legal briefs and the case law they printed from the legal information services to do their legal research; we used this material informally at the time to verify what the students had said in interviews about their annotation practices. In this study, we focus in greater detail on these extensive materials to better understand how the students' annotations contributed to the text of the briefs they wrote for Moot Court. The next section presents the study method and identifies specific types of annotation marks and their frequency near passages later used in the written brief. Based on these patterns, we have designed and developed a mechanism to identify, categorize, group, and rank annotations. We describe this mechanism, the mark parser, which was integrated into XLibris [8], and discuss how it can support readers in using their high-value annotations in subsequent writing tasks.

2 Relationship Between Annotation and Use

Legal briefs are written arguments that present one side's point of view in a legal case within the context of legal precedent. In the US, legal precedent is set by cases, published accounts of the facts and outcomes of legal rulings. Thus, students' writings refer to many prior cases — pointing out the similarities to and differences from the current case — to bolster their arguments. These references are in the form of quotations and citations.

Analysis of the students' legal briefs and the corresponding annotated court cases can yield insight into how different types of annotation were used to identify potential quotations and citations. Students used their annotations in other ways, such as reusing the structure of an argument, but in this analysis we focus on the connection between annotated passages, and quotations and citations used in the briefs. Reuse of references to annotated passages in writing is a conservative estimate of the utility of the annotations.

We spent a three month period observing the complete Moot Court process, from the initial period of instruction through the students' research and writing to the final mock courtroom presentation of competing oral arguments. During this period of time, we conducted a series of in-depth semi-structured contextual interviews of nine of the students who participated in the competition and the two faculty members and two third year law students who organized it. After the competition had ended, we collected the students' completed legal briefs and the printed-out cases they had read and annotated while they were doing their research. In [5] we mainly relied on field notes from the observations, full transcriptions from the interviews, and a set of digital photos to report important qualitative findings about legal research as it is conducted in this kind of heterogeneous information environment; we also describe how we used these findings to redesign an evolving tablet computer-based e-book prototype, XLib-

ris. The earlier study thus forms a backdrop for the more detailed investigation of the particular aspect of the law students' practice we discuss here: how readers – the law students – used their annotated source materials – the printed cases – in a writing task, composing a legal brief.

The study described in this paper is based on materials collected from seven students. The documents include each student's final brief and the printed cases the student read in preparing the brief. Each student wrote approximately 20 double-spaced pages and used from 46 to 105 legal references/citations. In addition to gathering and printing cases from the two online legal information services, Westlaw and Lexis, the students used textbooks, books in the library, and read online. As a consequence not all of the cases the students read and cited were available for our study. Thus our analysis is limited to the collection of those cases printed out on loose paper. According to the data gathered in our earlier interviews, these were generally the most recent cases and, for most students, were the most highly cited and relevant cases. Table 1 summarizes the number of documents collected, the number of citations in the student's brief, and the number of those citations that were to collected documents.

Table 1. Number of collected source documents and the percentage of student citations to them.

		Brief 1	Brief 2	Brief 3	Brief 4	Brief 5	Brief 6	Brief 7
A	# of collected source documents	16	11	20	13	21	10	27
B	Total # of citations	54	59	46	46	105	67	99
C	Citations to collected source documents	36 67%	45 76%	32 70%	46 100%	80 76%	23 34%	94 95%

2.1 From Citation to Source

The first detailed analysis of the documents identified the types of marks associated with passages that were subsequently used in the writing of the legal brief. In this analysis we looked at all citations in a brief and went back to locate the cited passage to see if and how it was annotated. For example, one brief included the statement:

“The court in *Vernonia* stated that the 'most significant element' of the case was that the drug testing program 'was undertaken in furtherance of the government's responsibilities, under a public school system, as guardian and tutor of children entrusted to its care.' *Vernonia*, 515 U.S. at 664.”

This citation comes from the annotated passage shown in Figure 1. We take a passage to be the identifiable region with which an annotation is associated.

Table 2 shows the number of citations in each brief that refer to documents collected in the study and the number of those citations where the passage in the original document was marked in a particular way. Places where a cited passage was marked in more than one way, such as with highlighter and a comment, have been included in the value for all applicable rows (e.g. highlight, comment, and multiple mark types). The

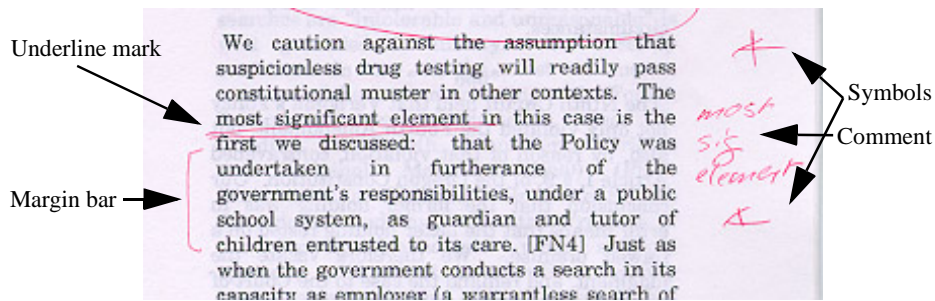


Figure 1. Annotations on section of case document.

rows starting with “highlight” are all subsets of the “marked in any way” set and indicate how many marks of that type were used. For example, the citation of the passage in Figure 1 is included in five counts for this brief: Marked in any way, Comment, Underline, Symbol, and Margin bar. These types of marks are based on opaque digital ink; highlights are similar to underlines using translucent ink and circles are looped text in either opaque or translucent ink — the bottom of a circle can be seen at the top of Figure 1.

Table 2. Number of citations and types of annotation on cited passages for each student.

	Brief 1	Brief 2	Brief 3	Brief 4	Brief 5	Brief 6	Brief 7
Citations to collected source documents (row C in Table 1)	36	45	32	46	80	23	94
Unmarked	8	8	27	5	78	10	26
Marked in any way	28	37	5	41	2	13	68
Highlight	27	1	0	35	0	10	43
Comment	3	14	0	11	1	0	23
Symbol	8	4	0	7	1	0	5
Underline	1	5	4	1	0	2	5
Circle	0	17	0	5	0	1	0
Margin bar	0	14	1	0	0	0	18
Multiple mark types	12	18	0	17	0	0	25

We can characterize the varying annotation styles of the seven students by looking at their papers and by the data in Table 2 above. The first brief writer used highlighters and symbols to code the documents as he read them. The second brief writer did not use highlighter, but instead used circles or margin bars and wrote comments in the margins of the documents. The third writer made few marks — 22 total passages across 20 cases. The fourth writer frequently combined highlighting and comments in the margin for areas of particular interest. The fifth writer made very few marks on the

papers: four symbols and one comment (not actually about the document annotated but containing information from an on-line source pertinent to the case). The sixth writer used more references to older cases from casebooks than other students did, which resulted in relatively few citations to the printed documents. This student used the highlighter or similar selection-oriented markings (underlines, circles, or margin bars), but did not use symbols or put comments in the margins. The last brief writer annotated using highlighter and pen marks. These stylistic characterizations confirm the data from student interviews in which they described their general annotation practices (see [5]).

From our review of the data collected in this study and from other related studies of annotation [4, 6, 9], we have observed clear stylistic differences in how people annotate. Some people prefer to highlight; some underline or circle; and some use margin bars, comments, and symbols. Despite all these differences, we can distinguish two general categories of annotators: those for whom annotation is a common practice (the “happy highlighters,” students 1, 2, 4, and 7), and those for whom it is the exception (the “meager markers,” students 3 and 5). Note that student 6 falls in between these two categories; in practice the two categories are actually stereotypes at the extremes of a continuum.

2.2 From Source to Citation

We now look at how frequently the students created single and multiple annotations for a passage. Table 3 summarizes readers' annotations of the source materials.

Table 3. Documents, pages, passages marked, and passages multimarked per student.

	Brief 1	Brief 2	Brief 3	Brief 4	Brief 5	Brief 6	Brief 7
# source documents (row A in Table 1)	16	11	20	13	21	10	27
# marked documents	15 <i>94%</i>	11 <i>100%</i>	2 <i>10%</i>	13 <i>100%</i>	2 <i>10%</i>	7 <i>70%</i>	22 <i>81%</i>
total # of pages in marked documents	279	178	36	331	78	231	608
# pages with marks	148 <i>53%</i>	98 <i>55%</i>	8 <i>22%</i>	102 <i>31%</i>	3 <i>4%</i>	69 <i>30%</i>	219 <i>36%</i>
# marked passages	552	325	22	311	5	159	688
passages marked per page marked	3.73	3.32	2.75	3.05	1.67	2.30	3.14
# multimarked passages	83	59	1	75	0	10	172
% of marked passages multimarked	15%	18%	5%	24%	0	6%	25%

The happy highlighters annotated (marked in any way) between 81 and 100% of their source documents; by contrast, two of the meager markers only annotated 10%

while student 6 annotated 70% of her source documents. Among the 1741 total pages of marked-up source documents, there were 647 marked pages with annotations on 2062 passages. Again, the happy highlighters marked a higher percentage of pages in their documents and had more passages marked per page than the meager markers. Where the happy highlighters really stood apart was in the use of multiple marks for a single passage. They had between 59 and 172 such passages in their source materials where the other students had between 0 and 10.

2.3 Discussion

One confounding factor in this analysis is we do not know how many times each document was read and whether the students annotated the documents during the first or a subsequent reading. While the annotations were still next to cited passages, it is possible the students remarked the passages while writing the brief; thus any mechanism we designed based on our analysis would apply only to re-reading or re-use of the materials for a similar task. However, during interviews with the students, they described re-reading key cases while they were outlining or writing. One student reported that she would go over her annotations during a second reading pass and put an asterisk by the important ones [5].

A second issue is whether the results from this specific annotation task can be generalized to inform the design of annotation systems. In particular, do people generally annotate heavily enough to imply value in algorithms for locating high-value annotations? While the writing task was different, O'Hara and Sellen's subjects describe annotation practices similar to those seen here [6].

3 Annotation Patterns and Useful Passages

In spite of the students' stylistic idiosyncrasies, this study indicates that there may be some simple methods for using mark grouping and type information to select passages that were relatively likely to be used in the students' writing task.

As a first pass, we look at the how selecting passages with any mark at all or just those passages that are multimarked might have helped these students. To do this we compare the measures of recall and precision for these two passage selection criteria. First, we define recall and precision measures based on percent of passages cited in the students' briefs that were annotated with any mark in the source materials. Second, we define recall and precision measures based on percent of passages cited in the students' briefs that were annotated in the source materials with multiple marks of different types, such as margin bar and comment, or circled passage and highlighted phrase. In the following discussion we refer to such sets of marks as multimarks. These measures were chosen to allow us to compare the relative utility of these criteria for sparse marking patterns, as well as for dense marking patterns (our so-called meager markers and happy highlighters). Table 4 summarizes these measures across all of the briefs.

By comparing the frequencies of cited passages that were annotated, we can evaluate different algorithms for locating useful passages based on annotation. For briefs 1,

Table 4. Recall and precision rates for passages with any mark and passages with multiple marks

	Happy Highlighters				Brief 6	Meager Markers	
	Brief 1	Brief 2	Brief 4	Brief 7		Brief 3	Brief 5
Mark recall	78%	82%	89%	72%	56%	16%	3%
Mark precision	5%	11%	13%	10%	8%	23%	20%
Multimark recall	33%	40%	37%	27%	0%	0%	0%
Multimark precision	14%	31%	23%	15%	0%	0%	--

2, 4, and 7 (the happy highlighters), 14%, 31%, 23%, and 15% of the passages with multiple marks were cited by students. This is the precision score for using multiple marks on a passage as a selection criterion. For students 3 and 5 (the meager markers), 23% and 20% of the passages marked in any way were later cited. It is evident that the happy highlighters and meager markers achieve comparable precision scores with multimarks and single marks, respectively. But student 6, who was in between the happy highlighters and the meager markers in annotation practice, ends up with only an 8% precision rate when using marks for retrieval and 0% when using multimarks.

Figure 2 presents another view of the precision and recall figures for cited passages in each brief. Precision percentage is calculated as the percent of marked or multimarked passages actually cited. The table shows that to maximize precision, we should use multimarks when they are sufficiently frequent (i.e. the ratio from Table 3 exceeds 0.1). In Table 4, bold indicates the initial method best suited for selecting useful passages. The choice of single vs. multiply marked passages makes these heuristics sensitive to the readers' predisposition to marking.




Figure 2. Precision and recall graph for selecting any mark (m) versus selecting multimarks (M) for the seven students.

Figure 2 shows the two clusters of students (“happy highlighters” and “meager markers”) with the selection methods performing significantly different for the two

groups. This clustering has led us to develop a simple annotation-style user model for determining the appropriate selection and emphasis criteria.

On one hand, the precision and recall rates are seemingly low, raising questions of their ultimate utility. These rates may be explained by the fact that the students' marks are not universally geared toward writing. On the other hand, detecting these patterns presents an opportunity to implement the annotation-based facilities we describe in this paper. It may be that these rates are sufficient to implement useful functionality, and will serve to reinforce this aspect of annotative practice.

In some cases, e.g. briefs 1, 6, and 7, our mark/multimark selection heuristic yields too many potentially important passages; and therefore is insufficiently selective for our purposes. Thus, while selecting multimarks first and then marks is likely to provide some benefit, further discrimination is required. One way to improve our heuristic is to develop a method of ranking the marked passages. We discuss such a method in the next section.

4 Grouping, Typing and Computing the Emphasis of Annotation Marks

Based on the analysis of annotation and citation patterns in legal briefs, we found a common set of base annotation types. There are two main categories of annotation types: selection marks such as highlighting, underlining, circling, and margin bars that identify a document passage, and interpretative marks such as comments and symbols that respond to the content. Of course, the roles of these two classes of marks are not completely distinct, since the act of selecting text out of a body is an interpretive act and marginalia naturally select text nearby as the focus of attention.

These observations led us to design and implement a mark parser that identifies structure and emphasis in pen stroke information of a reading appliance like XLibris [8]. The mark parser categorizes groups of marks into annotation types and constructs a passage-based grouping of annotations to identify high-emphasis sections of the document.

the Court's entanglement precedent to limit violative entanglement to institutional entanglement. *Jones I*, 930 F.2d at 423 (citing, *inter alia*, *Lynch*, 465 U.S. at 689, 104 S. Ct. at 1368 (O'Connor, J., concurring)). That a rabbi wrote and delivered the prayer at issue in *Lee* makes entanglement analysis relevant to that case, but the Resolution keeps Clear Creek free of all involvement with religious institutions. D. Endorsement

Like *Lemon's* advancement test, the Court's endorsement analysis focuses on the effect of a challenged governmental action. This is why, perhaps mistakenly, we conflated advertisement and endorsement analysis in *Jones I*. Because the Court has never tolerated a government endorsement of religion that is incidental to a primary secular effect, as it has with incidental religious advancements, we will not now compare

Figure 3. Multimarked text from XLibris.

To illustrate how the parser works, we will use the text and annotations from Figure 3. In this example our goal is to detect that the red ink marks are comments attached to the same passage as the neighboring highlighter marks. This would mark the two passages as multimarked and provide relative weightings of perceived emphasis (emphasis values) for comparing these and other marked passages.

The mark parser consists of a three-pass process of grouping, typing, and ranking the relative importance of reader marks. The first pass of the parser clusters the time-ordered marks into low-level groups based on time, distance, and pen type. This combines the individual pen strokes into clusters that are considered a single annotation.

The parser's first pass uses a function of temporal and spatial distance to determine breaks in the grouping of marks similar to that in Dynamite [2]. The algorithm includes pen-type differences (i.e., highlighter vs. ink) in computing the temporal and spatial gaps within the stroke sequence, since highlighters and pens are used differently in annotation. Practically, XLibris data demonstrates that temporal gaps between related marks (such as those strokes constituting a written annotation) are larger when using the pen than the highlighter. We speculate that readers may reflect more when they write interpretive comments and symbols than when they simply highlight a passage. The first stage of the parser also groups pen strokes that represent contiguous, but spatially distinct marks such as multi-line comments. This process results in a one or two level hierarchical clustering of marks into groups, as shown in the dotted and dashed regions in Figure 4. The dashed lines in Figure 4 indicate the results of this first stage. The dotted line regions are the results of the temporal and spatial clustering when multiple clusters are grouped together as being part of a single mark.

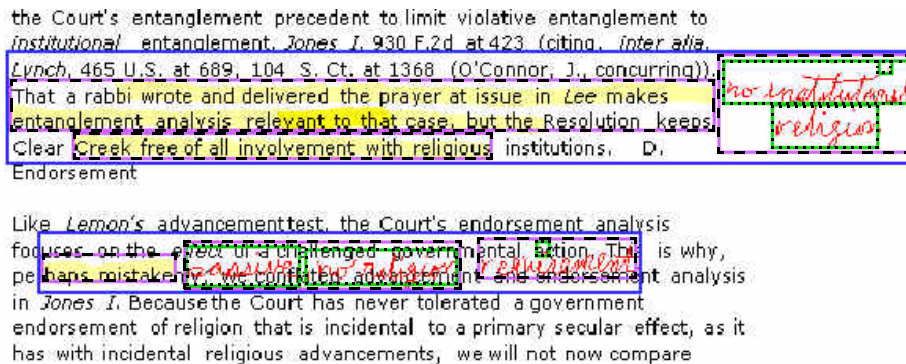


Figure 4. Clustering results are shown inside dotted-line rectangles (low-level clusters) and dashed-line rectangles (higher-level clusters). In the second pass, these groups are assigned types, attached to text passages, and grouped into higher-level structures (shown as solid lines).

Note that the parser broke up the top set of highlighter strokes due to a large temporal gap between the strokes; the interlinear comment (ink) on the bottom passage is similarly broken up. There is a trade-off between false-positive (grouped when they should not be) and false-negative (not grouped when they should be) when defining a boundary in spatial/temporal clustering. Because the later passes of the mark parser

combine clusters into higher-level structures, we have adopted a conservative strategy that produces more clusters. These smaller clusters are likely to be grouped in subsequent passes of the mark parser.

The second pass categorizes the mark clusters into annotation types using the pen type, the number of component marks, and the size and shape of the extent of the cluster. In the example in Figure 4, the six higher level clusters (shown in dashed-line rectangles) include three highlight annotations and three comment annotations.

Stroke types are based on pen type, the number of overlapping words, and the stroke's shape. For example, clusters composed of five or more pen marks (as opposed to highlighter marks) are considered comments. When the clusters are categorized, emphasis values are also assigned; the emphasis values are heuristic assessments of the relative importance of the annotation types.

There are three basic emphasis values assigned to the mark clusters. These values are based on two design guidelines. First, initial emphasis values were selected so that passages with a combination of mark clusters would be ranked higher than or equal to passages marked with only one cluster. Second, marks that are more precise in indicating focused attention are assigned higher emphasis values than those that are just rough indications of interest. So for the area highlight and area circle marks, the emphasis value is low (10). More focused selection marks — underline, margin bar, highlight, and circle — are assigned an intermediate value (15). Interpretive marks (comments, symbols, and callouts) are given a high emphasis value (20). In our example, the three comment clusters have an emphasis value of 20, while the three highlighter clusters have emphasis values of 15.

The third pass of the mark parser groups the mark clusters with passages from the source text. Text passages are associated with individual strokes by identifying the sentence nearest the stroke, with some preference for sentences next to the marks (thus a greater horizontal gap is allowed relative to the vertical gap). Any textual overlap between the passages associated with mark clusters will cause them to be grouped into a multimark cluster, creating the highest-level entities in the resulting parse tree. The highlighter and comments are attached to the same sentences in the text in the above example so the two passages are correctly recognized as multimarks, shown by solid rectangles, in Figure 4.

Emphasis values of these multimark clusters are the sum of the emphasis values of the component clusters. Thus, our two marked passages in the example above have emphasis values of 50 and 55 for the top and bottom passages respectively.

As the example illustrates, there is ample opportunity for misrecognition in the parsing process. Stroke clustering, attaching type information, and grouping passages with clusters are heuristic but the overall results match the desired outcome. Because lower levels of the parse may contain incorrect sub-groupings, it is important to design uses for the recognized structure that do not rely on precise clustering results.

5 Using the Results of the Mark Parser

The motivation for the mark parser was to locate areas of high emphasis that could help focus a reader's attention in re-reading a document or in performing a related task




Figure 5. Emphasis marks in XLibris's thumbnail overview of a document

such as writing. The mark parser has been integrated into XLibris to aid in locating valuable passages. XLibris readers often navigate to particular sections of a document using the thumbnail overview. By indicating annotation type and emphasis in this view, the user's attention will be directed to high-emphasis passages. Figure 5 shows the current use of color and icon shape to bring attention to particular passages in the document. This has similarities to Graham's depiction of computer-generated annotations in thumbnails [3] although XLibris' additions are based on the reader's own interactions with the document.

Color and icon are varied according to the emphasis value and type of the annotations. Color is used to indicate the relative emphasis of marked passages. Passages with high emphasis values (above 49) appear in red; those with medium values (between 35 and 49) appear in orange; and the others over the threshold (between 23 and 34) appear in yellow. The two icons used in the overview represent information about the type of the annotations on the emphasized passage. Asterisks are put next to annotations including interpretive marks such as comments, symbols, and callouts. An emphasized passage with only selection marks (highlighters, underlines, margin bars, and circles) is marked by a tab or exaggerated margin bar on the side of the page. Currently, we only use the two icons but this can be extended to convey more of the annotation type information to the user.

6 Summary

Our study of the law students' materials reinforced our past findings on the idiosyncratic nature of annotation practices. Within the variety of annotation styles, however, students could be grouped based on whether annotation is a common activity or whether it is the exception. On one hand, students who annotated sparsely were more likely to cite any passage they had annotated. On the other hand, those who annotated frequently were more likely to cite any passage with multiple marks.

A precision/recall analysis of using selection criteria based marked and multimarked passages indicated potential for such metrics but also pointed out the need for a more refined scheme for comparing marked passages.

This analysis led to the design of the mark parser. The mark parser uses location, time, and pen type information to cluster strokes into annotations. It then uses clusters' location, size, shape, pen information, and stroke information to categorize them into one of the following annotation types: comment, symbol, callout, highlight, circle, margin bar, or underline. Finally, the mark parser groups annotation clusters together based on the passage of the document to which the annotation refers. Relative emphasis values are associated with the annotation clusters and with the passage-based groups of annotations. By using the clustering and grouping heuristics in conjunction with the assessment of the relative value of the reader's annotations, we can identify which marks are likely to represent high-value annotations.

Initial experience running the mark parser on XLibris annotations gathered during an earlier study of a reading group [4] indicates that it correctly identifies the majority of high-level annotations, although lower-level structures in the parse include some noise due to the heuristic nature of the process. As an example of how the parse results can be used to help readers in their subsequent activities, we implemented a thumbnail-based document overview that indicates where the high-value annotations are and allows the reader to navigate directly to them. This application exemplifies how the mark parser results enable functionality and user interface improvements supporting reading-related activities, such as writing and collaboration, for systems supporting a combination of freeform and other annotations [7].

References

1. Adler, A., Gujar, A., Harrison, B.L., O'Hara, K., and Sellen, A. A Diary Study of Work-Related Reading: Design Implications for Digital Reading Devices, in Proceedings of CHI '98 (April 1998), ACM Press, 241-248.
2. Chiu, P. and Wilcox, L. A Dynamic Grouping Technique for Ink and Audio Notes, in Proceedings of UIST '98 (November 1998), ACM Press, 195-202.
3. Graham, J. The Reader's Helper: A Personalized Document Reading Environment, in Proceedings of CHI '99 (May 1999), ACM Press, 481-488.
4. Marshall, C.C., Price, M.N., Golovchinsky, G., and Schilit, B.N. Introducing a digital library reading appliance into a reading group, in Proceedings of ACM Digital Libraries '99 (August 1999), ACM Press, 77-84.
5. Marshall, C.C., Price, M.N., Golovchinsky, G., and Schilit, B.N. Designing e-Books for Legal Research, in Proceedings of JCDL 2001 (June 2001), ACM Press, 41-49.
6. O'Hara, K. and Sellen, A. A comparison of reading paper and on-line documents, in Proceedings of CHI '97 (March 1997), ACM Press, 335-342.
7. Phelps, T. and Wilensky, R. Multivalent Annotations, in Proceedings of the 1997 European Conference on Digital Libraries (September 1997), Springer, 287-303.
8. Schilit, B.N., Golovchinsky, G., and Price, M.N. Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations, in Proceedings of CHI '98 (April 1998), ACM Press, 249-256.
9. Wolfe, J.L. Effects of Annotations on Student Writers, in Proceedings of ACM Digital Libraries 2000 (June 2000), ACM Press, 19-26.