Hi-Cites:
Dynamically Created Citations with Active Highlighting

Michelle Q Wang Baldonado*
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304 USA
baldonad@parc.xerox.com

Terry Winograd
Gates 3B
Stanford University
Stanford, CA 94305-9035 USA
winograd@cs.stanford.edu

ABSTRACT
The original SenseMaker interface for information exploration [2] used tables to present heterogeneous document descriptions. In contrast, printed bibliographies and World Wide Web (WWW) search engines use formatted citations to convey this information. In this paper, we discuss *hi-cites*, a new interface construct developed for SenseMaker that combines the benefits of tables (which encourage the comparison of descriptions) and citations (which facilitate browsing).

Hi-cites are dynamically created citations with active highlighting. They are useful in environments where heterogeneous structured descriptions must be browsed and compared with ease. Examples beyond digital libraries include product catalogs, classified advertisements, and WWW search engines.

We have performed an evaluation of hi-cites, tables, and citations for tasks involving single attribute comparisons in the digital-library domain. This evaluation supports our claim that hi-cites are valuable for both comparison and skimming tasks in this environment.

Keywords
Hi-cites, dynamic citations, highlighting, digital libraries, information visualization, browsing

INTRODUCTION
Look for a book entitled “Physics” in most library card catalogs and you will be amazed at the number of books bearing that name. Fortunately, library catalogers describe documents by recording values for a wide variety of characteristic attributes. Library patrons can learn the author, publication year, number of pages, and the Library of Congress subject heading for each of the “Physics” books that have entries in a library’s card catalog.

Users who perform information exploration in digital libraries spend much time perusing document descriptions of the kind found in library card catalogs. However, the descriptions come from an ever-growing variety of sources. Information may come not only from on-line library catalogs but also from video archives, World Wide Web (WWW) search engines, music indices, e-mail archives, map databases, and much more. The structured descriptions obtained from these sources can be heterogeneous in two ways. First, these sources may employ different conventions for describing documents. For example, USMARC defines one set of descriptive attributes, while Z39.50 Bib-1 defines a different set. In the Stanford Digital Library Project, we have addressed this problem through the development of a metadata architecture [1] that accommodates different descriptive attribute sets as well as services for translating descriptions from one attribute set to another.

A second type of document-description heterogeneity revolves around document genre. The choice of attributes for a document’s description often depends on its genre. Document descriptions for music recordings are likely to include information about performers, tempo, and style; descriptions for WWW pages are likely to include information about URL and file size. The SenseMaker interface has been designed with this second style of heterogeneity in mind. The overall goal of SenseMaker is to support the contextual evolution of a user’s interests during information exploration in the digital library (see [2] for more details). This paper focuses solely on how to present the heterogeneous document descriptions obtained while interacting with SenseMaker.

TABLES AND CITATION SETS
The original SenseMaker interface presented search results (document descriptions) in a table whose columns were dynamically defined by users. Tables encourage users to compare descriptions in terms of a particular attribute, due to our visual disposition to use whitespace and proximity as perceptual grouping cues [5]. Focusing on a single attribute while browsing a collection allows users to gain an overview of the collection with respect to that attribute. For example, focusing on *publication location* while browsing a collection of results produced by a search for “ODA” might lead the user to observe that much work on ODA has taken place in Europe. This style of overview is particularly useful in heterogeneous environments where document descriptions come from widely varying contexts. In a pilot user study performed on the original SenseMaker interface (in which users gathered references for a hypothetical term

Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI ’97), Los Angeles, California, April 1998.

* When this paper was first submitted, the author was a Ph.D. student at Stanford University.

paper on a given topic), users explicitly mentioned this ability to see all values for a particular attribute at once as one of the features that they particularly liked.

Figure 1 shows an excerpt from the original table-based version of SenseMaker. Each row in the table corresponds to one document. Each column in the table stands for an individual attribute of the document.

Figure 1: Excerpt from Table-Based Version of SenseMaker

As this figure illustrates, a tabular view of heterogeneous document descriptions is often sparse. We have *URL* values but no *author* values for results from WWW search engines, and vice versa from citation databases. The multiplicity of "*Not available*" values is distracting to users. Furthermore, users must almost always use horizontal scrolling since screen space disappears quickly when each attribute requires its own column. Skimming is also difficult due to the spatial discontinuities that are characteristic of tables.

In contrast to tables, formatted bibliographies require less screen space and provide a spatially continuous flow of information. A bibliographic citation concatenates just the attribute values that are available for the document it describes. Furthermore, these citations make use of conventions that indicate for each value what its associated attribute is. These conventions may rely on fixed order (most bibliographic styles prescribe a particular order for citation elements) and may also use visual cues such as font style. A typical citation appears in Figure 2.

> Strunk, W. Jr. and White, E. B. *The Elements of Style.* Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Figure 2: Typical Citation

In this figure, italics are used for book titles, while the order of author, book title, edition, publisher, publication city, and publication year is conventional. The demarcating punctuation marks between items serve to delineate and identify each item in the citation as well.

In the publishing world, standard bibliographic styles are used for traditional written works, such as books, magazine articles, newspapers, memos and more. Formatted bibliographies can be generalized to accommodate nontraditional document descriptions as well (e.g., descriptions of videos, software, e-mail, maps, people, etc.). We call such a generalized bibliography a *citation set*. Note that most World Wide Web search engines present their results as citation sets (see Figure 3) .

Figure 3: AltaVista Citations

The gain in browsing ease offered by citation sets (due to the compactness and visually identifiable structure of citations) is offset by a loss in ease of comparison of available items. In this paper, we define and propose *hi-cites*, a novel interface idea that promises to improve browsing by combining the benefits of citation sets with the benefits of tables. The SenseMaker interface has been revised to make use of hi-cites instead of tables.

HI-CITES

Hi-cites, dynamically created citations with active highlighting, are a hybrid of citation sets and tables. They are like citation sets in that they concatenate visually marked attribute values (subject to wraparound), treating attribute values as though they were words and phrases in a single sentence. Definable citation styles encapsulate the ordering, marking, and rendering rules required for this display. For example, one citation style might stipulate that *title* appear before *author*, while another might change that ordering. In addition, one citation style might mark an *editor* value with the suffix "ed.," while another might use the suffix "(editor)." Finally, a citation style might specify that article *titles* are quoted while journal *titles* are italicized.

Hi-cites are like tables in that they facilitate the perceptual grouping of all values of a particular attribute. Specifically, users can highlight* in red all values for an attribute by pausing for a short period of time over that attribute with the mouse (see the next section for a detailed discussion of this affordance). Figure 4 shows both *title* highlighting and *publisher* highlighting. Note that it substitutes reverse video for color so that the highlighting can be represented in a black-and-white figure.

In addition to enabling highlighting, the point-plus-delay affordance also allows users to find out the name of each attribute. If a SenseMaker user were to move her cursor over a *title* attribute value and pause, then all *title* values would be highlighted. In addition, the name *Title* would

* The name hi-cite derives from the words "highlight" and "citation."

Strunk, W. Jr. and White, E. B. *The Elements of Style*. Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Morris, W. and Morris, M. *Morris Dictionary of Word and Phrase Origins*. Harper and Row, Publishers, New York, 1977.

Strunk, W. Jr. and White, E. B. *The Elements of Style*. Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Morris, W. and Morris, M. *Morris Dictionary of Word and Phrase Origins*. Harper and Row, Publishers, New York, 1977.

Figure 4: Two Sets of Hi-Cites (actual hi-cites make use of color rather than reverse video for highlighting)

show up in a temporary rectangular pop-up box. This result is shown in Figure 5.

This "help" feature is important for cases where users are not familiar with the chosen attribute marking conventions. For example, a newcomer to geographic information sources may not be familiar with the conventional marking of latitude and longitude, and thus may need to discover what the pieces of a geographic citation are.

Affordances

The hi-cite affordance for highlighting an attribute and for discovering an attribute's name (point-plus-delay over an individual attribute value) is similar to that needed to initiate "tool tips" in Windows applications. The choice of this affordance was made after comparing it to a number of other possibilities. In an early design that we considered, the affordance was simply moving the mouse over an attribute value, similar to the affordance for Macintosh "balloon help" [8]. However, we discovered that this was distracting to the user. Furthermore, the interaction between scrolling and skimming was problematic. Consider the situation where a user wants to skim all titles in the collection of document descriptions (not just those that fit into the current window). If moving the mouse over an attribute value causes highlighting to occur, then the act of moving the mouse to the scrollbar is likely to cause the desired highlighting to be lost.

We rejected the possibilities of highlighting on single-click or double-click on the grounds that they would conflict with the design languages [7] of Web browsers and GUIs, both of which are already well-known and familiar to our intended user community. In particular, single-click maps onto "follow link" in Web browsers, while double-click maps onto "open" in GUIs.

Finally, we debated making the affordance something less familiar, such as typing a letter over the attribute value (e.g., "h" for highlight); right-clicking and then either automatically highlighting or bringing up a pop-up menu (as in Windows); or doing a combination of special key (e.g., the Apple key or the Control key) and clicking. In all of these cases, the affordance would be learnable, but it was unlikely that users would "stumble" across the affordance. We felt that highlighting was so crucial to hi-cites that ease of discovery was as important of ease of learning. Thus, we decided to use the point-plus-delay affordance, even though the delay means that highlighting is not as fast as it could be for the user. Clearly, choice of delay time deserves further study.

Implementing Citation Styles

Many different citation styles exist in the publishing world. Hi-cites also allow for different citation styles. In SenseMaker, a citation style is implemented as an ordered list of attributes. In addition, objects specify how each attribute should be rendered (e.g., italicized or quoted) and marked (e.g., with a suffix "ed."). In our implementation, attribute values can be rendered as links (pointers) to the actual attribute values. This use of links in citations has already become common practice in the reference sections of documents on the World Wide Web. For example, articles in the on-line D-Lib magazine (http://www.dlib.org/) often include bibliographic citations that have hypertext links to the referenced works. Similarly, the result format for the various WWW search engines is an implicit citation style that contains links.

Including link information in the rendering-style object also allows local values that are particularly long to be accessible via a link rather than directly integrated into the citation. Consider for example the difference in expected lengths of a title, abstract, and full text of a document. Titles are usually fairly short in length and are almost always listed in on-line citations. Abstracts are usually on the order of one to two paragraphs. In many electronic card catalogs, the user has a choice of a terse or verbose presentation of citations, where the abstract is listed in only the verbose style. Finally, full text is almost never included directly in the citations of an on-line catalog, but may be available via a special command. The ability to have links in citations changes the design space in that it makes it possible to have longer values readily accessible without giving up screen space.

In non-textual domains, such as video, the "citation" for a video clip may be a multimedia sequence of images and text rather than an attribute-value style citation. In general, the question of how to display succinct representations of multimedia documents is an open area for research. We believe that at least some of the hi-cite principles are applicable to this domain. For example, the last hi-cite that appears in Figure 5 is a description of a video. It contains a *title* attribute, which is common to all of the displayed hi-cites, but it also contains a number of genre-specific attributes (e.g., *filmstrip*) that are accessible via links. An alternative to this citation style might inline the *filmstrip* value for each video directly into the hi-cite.

RELATED WORK

Highlighting

The use of highlighting as a mechanism for bringing a user's attention to a particular piece of information on the screen has been carefully studied by the human factors community. From this perspective, highlighting is defined as any visual characteristic that causes the highlighted piece of In addition to color, highlighting can include boxing, inverse video, blinking, and other techniques. Several studies have shown that color highlighting is useful for the visual-search task [3, 6]. Although no definitive answer has emerged on how to choose a color for highlighting, some studies show the perceptual salience of red [6]. In addition to choosing a color for highlighting, designers must also consider the amount of color in the interface.

One issue that has received great attention in highlighting studies is the effect of mis-highlighting information, or what happens when information is highlighted but the user does not wish to focus on it. In general, a high rate of mis-highlighting can override the benefits obtained from highlighting [6]. These findings are good evidence for the importance of user control in hi-cites. The fact that a hi-cite user can quickly and easily control what piece of information is highlighted means that the user can ensure that highlighting always matches the current focus of attention. These study results also indicate the need for SenseMaker users to be able to turn off highlighting completely in the interface, a feature that is not currently implemented.

One possible alternative to hi-cites is to differentiate the attributes within a citation by displaying each of them in a different color (e.g., red for *title*, blue for *publication date*). In this case, dynamic highlighting would not be necessary. Users could compare values for an attribute by focusing on the color corresponding to that attribute. However, this strategy requires addressing the issue of color blindness. Furthermore, we found informally in experimenting with prototypes that a multiplicity of colors in the interface was often distracting and detracted from our ability to perform perceptual grouping for a given attribute. This anecdotal evidence, combined with the literature on the effects of mis-highlighting, convinced us to design for dynamic, single-color highlighting rather than for static, multi-color highlighting.

Keyword-in-Context Indices

Another alternative to our hi-cite design originates in the pioneering work of H.P. Luhn on keyword-in-context indices (KWIC indices) [4]. A typical keyword-in-context index has one entry for each keyword of each indexed work's title. These titles are presented "in context," meaning that they are presented sandwiched in between the surrounding words from the title. The keyword entries themselves are aligned in a single column. In a classic keyword-in-context index, there was a limit on how many characters could be displayed. Therefore, a fixed number of characters from the prefix to the keyword could be included to its left and a fixed number of characters from the keyword to the suffix could be included to its right. An example of a keyword-in-context index appears in Figure 4.

In KWIC indices, position is used to highlight a piece of structured information. "The initial letters of the alphabetized keywords form a column which guides the eye when scanning for desired words" [4, p. 289]. In other words, the "highlighting" mechanism of keyword-in-context is alignment rather than color. This reliance on alignment means that tables are more similar to keyword-in-context indices than to hi-cites.

In a static display, highlighting by alignment is easier for scanning than highlighting by color since the eye does not need to rove across the page. However, dynamically changing what is highlighted is much more disruptive for highlighting by alignment than for highlighting by color since it causes the rearrangement of all the visible document descriptions.

In the heterogeneous digital-library domain, users are likely to change what is highlighted quite often. It is for this reason that we hypothesize that hi-cites are better suited for skimming document descriptions in this environment than are KWIC indices. Hi-cites keep document descriptions globally constant in space throughout the browsing process while still allowing perceptual grouping by color.

```
ve Survey of the Radio, Phonograph, and Television Habits in
te Manual of the Edison Phonograph
                    The Phonograph and its Inventor, Thomas
ete Library of American Phonograph Recordings
    The Edison Improved Phonograph
           The Fabulous Phonograph; from Edison to Stereo
e Patent History of the Phonograph, 1877-1912: A Source Book
                        Recollections of Thomas A. Edison: A
```

Figure 4: A Keyword-in-Context Bibliographic Index

EVALUATION

Hi-cites are modeled after both citation sets and tables. Citation sets are easier to skim than tables since they take up less screen space and provide a more spatially continuous flow of information. Hi-cites are so close in appearance to citation sets that it is reasonable to believe that they are also easier to skim than tables. Tables are better than citation sets for the comparison of attribute values since they facilitate perceptual grouping. Hi-cites are different enough from tables that an evaluation is necessary to determine if they are also better than citation sets for this task. Before conducting such an evaluation, we enumerated the following hypotheses about the differences we expected to find among tables, citation sets, and hi-cites for tasks that involve looking for particular attributes within a set of document descriptions or for performing single attribute comparisons of document descriptions.

- Task-completion times will be fastest for tables, then hi-cites, then citation sets
- Perceived task-completion times will be fastest for tables, then hi-cites, then citation sets
- Error rates will be smallest for tables, then hi-cites, then citation sets
- Hi-cites will be the preferred condition since they allow for rapid answering of the questions as well as for quick skimming.

With these hypotheses in place, we constructed an experimental design that allowed us to determine their validity.

Experimental Design

Fourteen subjects participated in our evaluation of hi-cites. All subjects were from Stanford University and were paid for their time. Thirteen of the subjects were either undergraduates or graduate students. One subject was a Stanford staff member in the Computer Science Department. No subjects were students in the Computer Science Department, although the majority were students in the sciences.

The study compared three different presentation styles: tables, citation sets, and hi-cites. Each subject completed a practice task (where textual and oral help were provided), and then proceeded to complete two actual tasks (no help provided). Each task had three steps, where each step involved viewing a new collection of document descriptions in a particular presentation style and then answering questions about those descriptions. Step order and the pairings between collections and presentation styles were varied randomly. The questions asked in each step were of the following types (and were always presented in this fixed order):

- How many descriptions have the word X in the title?
- How many descriptions have publication locations listed?
- How many descriptions have the publication date Y?

All of these questions require that a user consider the document collection from the perspective of a particular attribute. We opted for providing users with this standard set of questions rather than providing a real-world task because we felt that these questions ensured that users would indeed perform attribute-based comparisons and obtain attribute-based overviews of the document collection (the focus of our evaluation, as well as a useful strategy in real digital-library tasks). Also, these question-based tasks were amenable to statistical comparison.

The time to complete each step and the error rate for each step were measured automatically. At the end of each task, subjects were also asked to rank each presentation style in terms of perceived speed. Finally, at the completion of the study, subjects were asked to compare/contrast orally the different styles and to choose one as preferred.

Experimental Results

Completion times

The comparison of completion times for each condition used the average of each subject's Task 1 time and Task 2 time. This average compensates for potential differences in the ease of learning each step, since step order was varied randomly for each subject. A one-way repeated-measures ANOVA on these average completion times ($F = 7.52$, $p = .003$) showed that the three conditions differed significantly at the .005 significance level. At significance level .05, the *post hoc* Student-Newman-Keuls test showed that the average completion times for tables and hi-cites were not significantly different. However, the average completion times for tables and citation sets and for hi-cites and citation sets were significantly different. Thus, tables and hi-cites are significantly faster for these tasks than are citation sets. The completion times for each condition, averaged across all subjects, are presented in Table 1.

Table 1: Average Completion Times (in seconds)

Tables	Citation Sets	Hi-Cites
43.081	51.969	45.477

Error rates

A one-way repeated-measures ANOVA on the error rates (again, averaged over the two tasks) showed no significant difference among the error rates ($F = .24$, $p = .786$). The error rates for each condition, averaged across all subjects, are presented in Table 2.

Table 2: Average Error Rates (number of wrong answers ÷ 3)

Tables	Citation Sets	Hi-Cites
0.10	0.13	0.11

Rankings of perceived completion times

In addition to looking at the actual differences in average completion times for each of the conditions, we also looked

at perceived differences in completion times. Users' perceptions of time do not always match up to real time. Furthermore, a user's perception that an interface is fast and facile can make a crucial difference in whether or not an interface will be used, no matter what the actual performance metrics might be for the interface.

A one-way repeated-measures ANOVA* ($F = 10.72$, $p < .005$) showed that the means for the average rankings (1 is fastest; 3 is slowest) were not equal. At a significance level of .01, the Student-Newman-Keuls test showed that all differences were significant. Hi-cites were perceived to be faster than tables, and tables were perceived to be faster than citation sets. The average rankings for each condition are shown in Table 3.

Table 3: Average Rankings of Perceived Completion Times (1 is fastest; 3 is slowest)

Tables	Citation Sets	Hi-Cites
1.82	2.68	1.50

Preferences

Subjects specified their overall preferred presentation style at the end of the session, not at the end of each task. This data was analyzed using a t-test, where the null hypothesis was that the probability of a subject choosing a condition as preferred was 1/3. The observed proportions of preference assignments for each condition appear in Table 4. Computing t for each condition showed that the observed proportions of preference assignments for hi-cites and citation sets were significantly different from 1/3 at the .05 level, but that the observed proportion for tables was not significantly different from 1/3 at the .05 level.

Table 4: Observed Proportions of Preference Assignments

Tables	Citation Sets	Hi-Cites
4/14	0/14	10/14

User feedback

At the end of a session, each subject was asked not only to select a presentation style as preferred, but also to compare and contrast all of the presentation styles. This feedback served to contextualize the preference choices with a concrete sense of why subjects found their preferred style to be useful and also gave a good indication of what was easy or hard to do in the other presentation styles. Overwhelmingly, subjects were enthusiastic about the value of hi-cites. They particularly liked the way color guided their eye to the "right place" in the collection. However, they found much to like about tables also. They pointed out that tables arranged the information in a format that would be very convenient to browse if it all fit on the screen. The need to scroll horizontally in tables was the most frequent complaint about their use. Subjects explained their difficulties with citation sets in terms of problems with honing in on the right piece of information in a description. They felt that color and space served to mark attributes well, but that slight changes in format and relative position were not good perceptual markers. A sampling of subject comments appears in Table 5.

Table 5: Sampling of Subject Comments

Hi-cites	caught my eye; much easier to go by color than location
	really great; focus on particular area; red drew visual attention
	with a lot of text, it just hops out at you
	stuck out compared to other ones
	don't know where to find first one (or there might not be a location)
Tables	hard to use because you need to scroll
	liked having it all mapped out
	nice; knew where to look
	maybe easier except you can't see it all on the screen
	everything structured in such a way you could do the same thing easily
Citation Sets	confusing; hard to pick out information being asked for
	more time to sort out information
	didn't really like them; just slower to get through
	tedious to read through words to look for information
	had to look through everything

* Technically, the use of the ANOVA is problematic here since the observations are not independent. However, the ANOVA is still the most useful statistic for this situation to the best of our knowledge.

Discussion

Before conducting the experiment, we set forth the following hypotheses about tables, citation sets, and hi-cites.

- Task-completion times will be fastest for tables, then hi-cites, then citation sets
- Perceived task-completion times will be fastest for tables, then hi-cites, then citation sets
- Error rates will be smallest for tables, then hi-cites, then citation sets
- Hi-cites will be the preferred condition since they allow for rapid answering of the questions as well as for quick skimming.

Statistical analysis of our experimental results leads to the following conclusions:.

- Task-completion times are faster for tables and hi-cites than for citation sets, but there is no significant difference between tables and hi-cites
- Perceived task-completion times are faster for hi-cites, then tables, then citation sets
- Error rates do not differ significantly among tables, citation sets, and hi-cites
- Hi-cites are the preferred condition

Our experimental results were somewhat surprising in that hi-cites either outranked tables or were not sufficiently different from tables in all cases. Thus, we can conclude that hi-cites are indeed a hybrid between tables and citation sets in that they:

- allow for ease of comparison
- allow for ease of skimming

Ease of comparison is demonstrated by our experimental results. Ease of skimming stems from the fact that hi-cites and citation sets are both compact representations of information that take up less screen space and are more spatially continuous than tables.

CONCLUSION

We have proposed the concept of hi-cites: dynamically created citations with active highlighting. The generation of hi-cites relies on access to attribute-value encodings of documents. The appearance of hi-cites is governed by citation styles that dictate what attributes should be presented, in what order they should be listed, and how they should be rendered and marked. Users control the active highlighting of hi-cites through a “tool tip” affordance. Point-plus-delay on any attribute value causes all values for that attribute to be highlighted in the citation set. We have chosen red color highlighting for hi-cites based upon our understanding of previous highlighting studies.

We have argued that hi-cites combine the benefits of both bibliographic citations and tables. The evaluation that we have performed supports our claim for a particular set of tasks and conditions. More insight into hi-cites could be gained by devising experiments to compare hi-cites to KWIC indices, to measure the ease of discovery for the point-plus-delay affordance currently used for hi-cites, and to determine the effects of different styles of highlighting (e.g., reverse video vs. red) in this context.

This paper has focused on the role of hi-cites in SenseMaker, an interface to a heterogeneous digital library. We believe that hi-cites can be useful in other environments that require heterogeneous structured descriptions to be browsed and compared with ease. Examples beyond digital libraries include product catalogs, classified advertisements, and World Wide Web search engines.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under Cooperative Agreement IRI-9411306. Funding for this agreement is also provided by DARPA, NASA, and the industrial partners of the Stanford Digital Libraries Project.

We would like to thank Cliff Nass and Vicky Reich for their valuable advice on experimental design and statistical analysis.

REFERENCES

1. Baldonado, M., Chang, C. K., Gravano, L., and Paepcke, A. Metadata for digital libraries: architecture and design rationale, in *Proceedings of DL ‘97* (Philadelphia PA, July 1997) ACM Press, 47-56.

2. Baldonado, M.Q.W., and Winograd, T. SenseMaker: An information-exploration interface supporting the contextual evolution of a user’s interests, in *Proceedings of CHI ‘97 (*Atlanta GA, March 1997) ACM Press, 11-18.

3. Brown, T.J. Visual display highlighting and information extraction, in *Proceedings of the Human Factors Society 35th Annual Meeting* (1991), 1427-1431.

4. Luhn, H.P. Keyword-in-context index for technical literature (KWIC Index). *American Documentation* XI, 4 (1960), 288-295.

5. Nielsen, J. *Usability Engineering*. Academic Press, Boston, MA, 1993.

6. Philipsen, G. Effects of six different highlighting modes on visual search performance in menu options. *International Journal of Human-Computer Interaction* 6, 3 (1994), 319-334.

7. Rheinfrank, J., and Evenson, S. Design languages, in T. Winograd (ed.), *Bringing Design to Software*. Addison-Wesley, Reading, MA, 1996.

8. Sellen, A., and Nicol, A. Building user-centered online help, in B. Laurel (ed.), *The Art of Human-Computer Interface Design*. Addison-Wesley, Reading, MA, 1990.

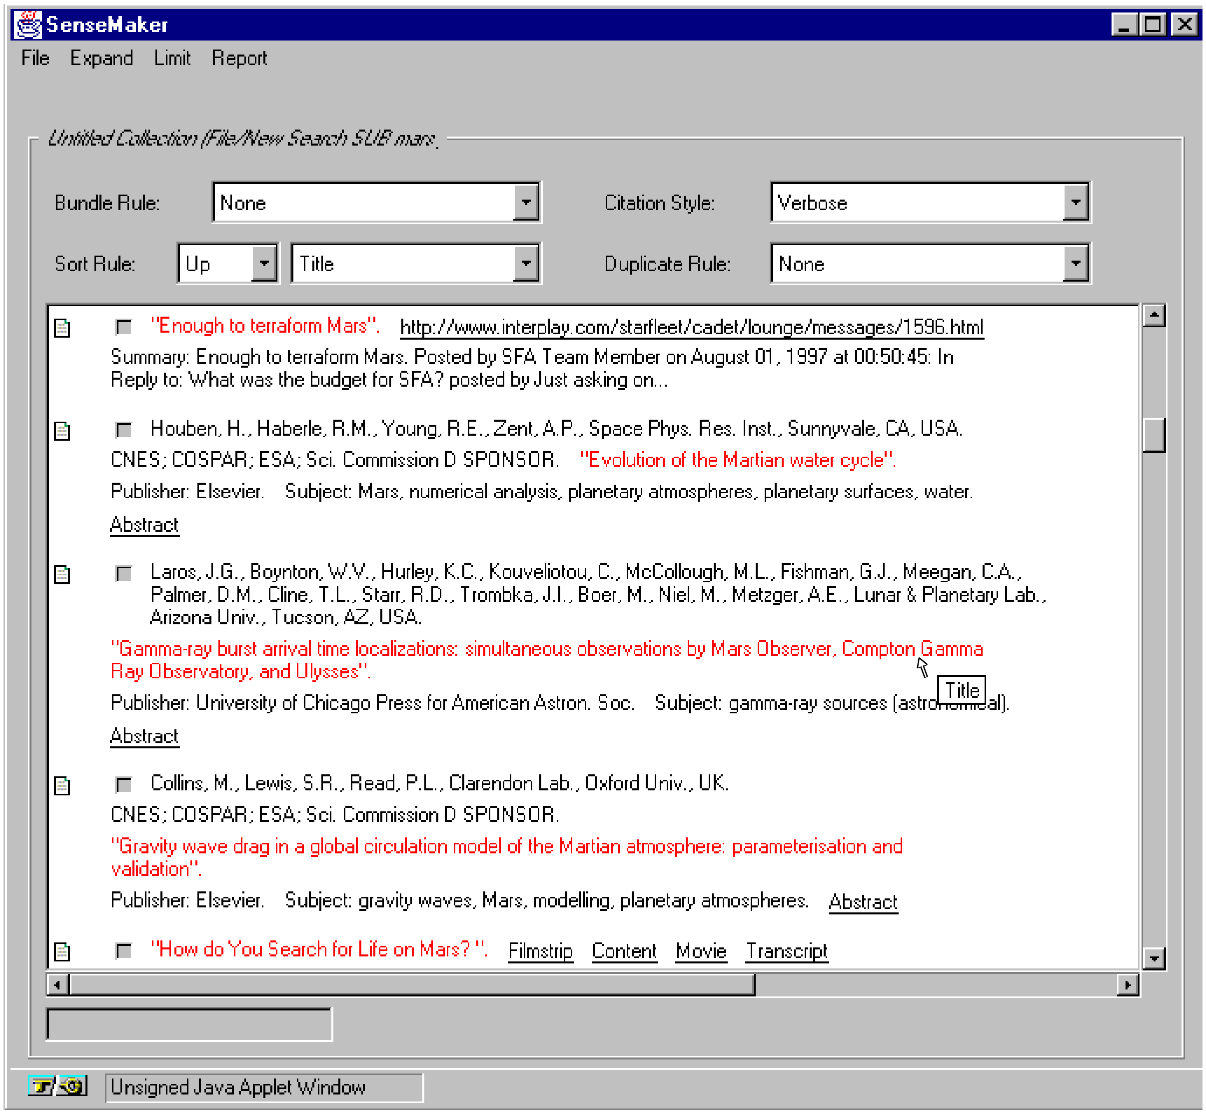

Figure 5: Hi-Cite Version of SenseMaker (*titles* are highlighted in red)
The heterogeneous set of documents described includes articles, WWW pages, and videos.

Hi-Cites:
Dynamically Created Citations with Active Highlighting

Michelle Q Wang Baldonado*
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304 USA
baldonad@parc.xerox.com

Terry Winograd
Gates 3B
Stanford University
Stanford, CA 94305-9035 USA
winograd@cs.stanford.edu

ABSTRACT

The original SenseMaker interface for information exploration [2] used tables to present heterogeneous document descriptions. In contrast, printed bibliographies and World Wide Web (WWW) search engines use formatted citations to convey this information. In this paper, we discuss *hi-cites*, a new interface construct developed for SenseMaker that combines the benefits of tables (which encourage the comparison of descriptions) and citations (which facilitate browsing).

Hi-cites are dynamically created citations with active highlighting. They are useful in environments where heterogeneous structured descriptions must be browsed and compared with ease. Examples beyond digital libraries include product catalogs, classified advertisements, and WWW search engines.

We have performed an evaluation of hi-cites, tables, and citations for tasks involving single attribute comparisons in the digital-library domain. This evaluation supports our claim that hi-cites are valuable for both comparison and skimming tasks in this environment.

Keywords

Hi-cites, dynamic citations, highlighting, digital libraries, information visualization, browsing

INTRODUCTION

Look for a book entitled "Physics" in most library card catalogs and you will be amazed at the number of books bearing that name. Fortunately, library catalogers describe documents by recording values for a wide variety of characteristic attributes. Library patrons can learn the author, publication year, number of pages, and the Library of Congress subject heading for each of the "Physics" books that have entries in a library's card catalog.

Users who perform information exploration in digital libraries spend much time perusing document descriptions of the kind found in library card catalogs. However, the descriptions come from an ever-growing variety of sources. Information may come not only from on-line library catalogs but also from video archives, World Wide Web (WWW) search engines, music indices, e-mail archives, map databases, and much more. The structured descriptions obtained from these sources can be heterogeneous in two ways. First, these sources may employ different conventions for describing documents. For example, USMARC defines one set of descriptive attributes, while Z39.50 Bib-1 defines a different set. In the Stanford Digital Library Project, we have addressed this problem through the development of a metadata architecture [1] that accommodates different descriptive attribute sets as well as services for translating descriptions from one attribute set to another.

A second type of document-description heterogeneity revolves around document genre. The choice of attributes for a document's description often depends on its genre. Document descriptions for music recordings are likely to include information about performers, tempo, and style; descriptions for WWW pages are likely to include information about URL and file size. The SenseMaker interface has been designed with this second style of heterogeneity in mind. The overall goal of SenseMaker is to support the contextual evolution of a user's interests during information exploration in the digital library (see [2] for more details). This paper focuses solely on how to present the heterogeneous document descriptions obtained while interacting with SenseMaker.

TABLES AND CITATION SETS

The original SenseMaker interface presented search results (document descriptions) in a table whose columns were dynamically defined by users. Tables encourage users to compare descriptions in terms of a particular attribute, due to our visual disposition to use whitespace and proximity as perceptual grouping cues [5]. Focusing on a single attribute while browsing a collection allows users to gain an overview of the collection with respect to that attribute. For example, focusing on *publication location* while browsing a collection of results produced by a search for "ODA" might lead the user to observe that much work on ODA has taken place in Europe. This style of overview is particularly useful in heterogeneous environments where document descriptions come from widely varying contexts. In a pilot user study performed on the original SenseMaker interface (in which users gathered references for a hypothetical term

Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '97), Los Angeles, California, April 1998.

* When this paper was first submitted, the author was a Ph.D. student at Stanford University.

paper on a given topic), users explicitly mentioned this ability to see all values for a particular attribute at once as one of the features that they particularly liked.

Figure 1 shows an excerpt from the original table-based version of SenseMaker. Each row in the table corresponds to one document. Each column in the table stands for an individual attribute of the document.

	Title	Author	URL
	Medical Image Volume Visualization Software FAQ	Not available	http://www.dca.fee.unicamp.br/.../medimage-FAQ.html
	Multimodality visualization of medical volume data	Zuiderveld, K.J.	Not available
	NPAC Technical Report SCCS-508	Not available	http://www.npac.syr.edu/.../sccs-0508.html
	Nicer-Slicer-Dicer (NSD)	Not available	http://mayflower.ncsa.uiuc.edu/.../nsd.html

Figure 1: Excerpt from Table-Based Version of SenseMaker

As this figure illustrates, a tabular view of heterogeneous document descriptions is often sparse. We have *URL* values but no *author* values for results from WWW search engines, and vice versa from citation databases. The multiplicity of "*Not available*" values is distracting to users. Furthermore, users must almost always use horizontal scrolling since screen space disappears quickly when each attribute requires its own column. Skimming is also difficult due to the spatial discontinuities that are characteristic of tables.

In contrast to tables, formatted bibliographies require less screen space and provide a spatially continuous flow of information. A bibliographic citation concatenates just the attribute values that are available for the document it describes. Furthermore, these citations make use of conventions that indicate for each value what its associated attribute is. These conventions may rely on fixed order (most bibliographic styles prescribe a particular order for citation elements) and may also use visual cues such as font style. A typical citation appears in Figure 2.

Strunk, W. Jr. and White, E. B. *The Elements of Style*. Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Figure 2: Typical Citation

In this figure, italics are used for book titles, while the order of author, book title, edition, publisher, publication city, and publication year is conventional. The demarcating punctuation marks between items serve to delineate and identify each item in the citation as well.

In the publishing world, standard bibliographic styles are used for traditional written works, such as books, magazine articles, newspapers, memos and more. Formatted bibliographies can be generalized to accommodate nontraditional document descriptions as well (e.g., descriptions of videos, software, e-mail, maps, people, etc.). We call such a generalized bibliography a *citation set*. Note that most World Wide Web search engines present their results as citation sets (see Figure 3) .

Figure 3: AltaVista Citations

The gain in browsing ease offered by citation sets (due to the compactness and visually identifiable structure of citations) is offset by a loss in ease of comparison of available items. In this paper, we define and propose *hi-cites*, a novel interface idea that promises to improve browsing by combining the benefits of citation sets with the benefits of tables. The SenseMaker interface has been revised to make use of hi-cites instead of tables.

HI-CITES

Hi-cites, dynamically created citations with active highlighting, are a hybrid of citation sets and tables. They are like citation sets in that they concatenate visually marked attribute values (subject to wraparound), treating attribute values as though they were words and phrases in a single sentence. Definable citation styles encapsulate the ordering, marking, and rendering rules required for this display. For example, one citation style might stipulate that *title* appear before *author*, while another might change that ordering. In addition, one citation style might mark an *editor* value with the suffix "ed.," while another might use the suffix "(editor)." Finally, a citation style might specify that article *titles* are quoted while journal *titles* are italicized.

Hi-cites are like tables in that they facilitate the perceptual grouping of all values of a particular attribute. Specifically, users can highlight* in red all values for an attribute by pausing for a short period of time over that attribute with the mouse (see the next section for a detailed discussion of this affordance). Figure 4 shows both *title* highlighting and *publisher* highlighting. Note that it substitutes reverse video for color so that the highlighting can be represented in a black-and-white figure.

In addition to enabling highlighting, the point-plus-delay affordance also allows users to find out the name of each attribute. If a SenseMaker user were to move her cursor over a *title* attribute value and pause, then all *title* values would be highlighted. In addition, the name *Title* would

* The name hi-cite derives from the words "highlight" and "citation."

Strunk, W. Jr. and White, E. B. *The Elements of Style*. Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Morris, W. and Morris, M. *Morris Dictionary of Word and Phrase Origins*. Harper and Row, Publishers, New York, 1977.

Strunk, W. Jr. and White, E. B. *The Elements of Style*. Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Morris, W. and Morris, M. *Morris Dictionary of Word and Phrase Origins*. Harper and Row, Publishers, New York, 1977.

Figure 4: Two Sets of Hi-Cites (actual hi-cites make use of color rather than reverse video for highlighting)

show up in a temporary rectangular pop-up box. This result is shown in Figure 5.

This "help" feature is important for cases where users are not familiar with the chosen attribute marking conventions. For example, a newcomer to geographic information sources may not be familiar with the conventional marking of latitude and longitude, and thus may need to discover what the pieces of a geographic citation are.

Affordances

The hi-cite affordance for highlighting an attribute and for discovering an attribute's name (point-plus-delay over an individual attribute value) is similar to that needed to initiate "tool tips" in Windows applications. The choice of this affordance was made after comparing it to a number of other possibilities. In an early design that we considered, the affordance was simply moving the mouse over an attribute value, similar to the affordance for Macintosh "balloon help" [8]. However, we discovered that this was distracting to the user. Furthermore, the interaction between scrolling and skimming was problematic. Consider the situation where a user wants to skim all titles in the collection of document descriptions (not just those that fit into the current window). If moving the mouse over an attribute value causes highlighting to occur, then the act of moving the mouse to the scrollbar is likely to cause the desired highlighting to be lost.

We rejected the possibilities of highlighting on single-click or double-click on the grounds that they would conflict with the design languages [7] of Web browsers and GUIs, both of which are already well-known and familiar to our intended user community. In particular, single-click maps onto "follow link" in Web browsers, while double-click maps onto "open" in GUIs.

Finally, we debated making the affordance something less familiar, such as typing a letter over the attribute value (e.g., "h" for highlight); right-clicking and then either automatically highlighting or bringing up a pop-up menu (as in Windows); or doing a combination of special key (e.g., the Apple key or the Control key) and clicking. In all of these cases, the affordance would be learnable, but it was unlikely that users would "stumble" across the affordance. We felt that highlighting was so crucial to hi-cites that ease of discovery was as important of ease of learning. Thus, we decided to use the point-plus-delay affordance, even though the delay means that highlighting is not as fast as it could be for the user. Clearly, choice of delay time deserves further study.

Implementing Citation Styles

Many different citation styles exist in the publishing world. Hi-cites also allow for different citation styles. In SenseMaker, a citation style is implemented as an ordered list of attributes. In addition, objects specify how each attribute should be rendered (e.g., italicized or quoted) and marked (e.g., with a suffix "ed."). In our implementation, attribute values can be rendered as links (pointers) to the actual attribute values. This use of links in citations has already become common practice in the reference sections of documents on the World Wide Web. For example, articles in the on-line D-Lib magazine (http://www.dlib.org/) often include bibliographic citations that have hypertext links to the referenced works. Similarly, the result format for the various WWW search engines is an implicit citation style that contains links.

Including link information in the rendering-style object also allows local values that are particularly long to be accessible via a link rather than directly integrated into the citation. Consider for example the difference in expected lengths of a title, abstract, and full text of a document. Titles are usually fairly short in length and are almost always listed in on-line citations. Abstracts are usually on the order of one to two paragraphs. In many electronic card catalogs, the user has a choice of a terse or verbose presentation of citations, where the abstract is listed in only the verbose style. Finally, full text is almost never included directly in the citations of an on-line catalog, but may be available via a special command. The ability to have links in citations changes the design space in that it makes it possible to have longer values readily accessible without giving up screen space.

In non-textual domains, such as video, the "citation" for a video clip may be a multimedia sequence of images and text rather than an attribute-value style citation. In general, the question of how to display succinct representations of multimedia documents is an open area for research. We believe that at least some of the hi-cite principles are applicable to this domain. For example, the last hi-cite that appears in Figure 5 is a description of a video. It contains a *title* attribute, which is common to all of the displayed hi-cites, but it also contains a number of genre-specific attributes (e.g., *filmstrip*) that are accessible via links. An alternative to this citation style might inline the *filmstrip* value for each video directly into the hi-cite.

RELATED WORK

Highlighting

The use of highlighting as a mechanism for bringing a user's attention to a particular piece of information on the screen has been carefully studied by the human factors community. From this perspective, highlighting is defined as any visual characteristic that causes the highlighted piece of In addition to color, highlighting can include boxing, inverse video, blinking, and other techniques. Several studies have shown that color highlighting is useful for the visual-search task [3, 6]. Although no definitive answer has emerged on how to choose a color for highlighting, some studies show the perceptual salience of red [6]. In addition to choosing a color for highlighting, designers must also consider the amount of color in the interface.

One issue that has received great attention in highlighting studies is the effect of mis-highlighting information, or what happens when information is highlighted but the user does not wish to focus on it. In general, a high rate of mis-highlighting can override the benefits obtained from highlighting [6]. These findings are good evidence for the importance of user control in hi-cites. The fact that a hi-cite user can quickly and easily control what piece of information is highlighted means that the user can ensure that highlighting always matches the current focus of attention. These study results also indicate the need for SenseMaker users to be able to turn off highlighting completely in the interface, a feature that is not currently implemented.

One possible alternative to hi-cites is to differentiate the attributes within a citation by displaying each of them in a different color (e.g., red for *title*, blue for *publication date*). In this case, dynamic highlighting would not be necessary. Users could compare values for an attribute by focusing on the color corresponding to that attribute. However, this strategy requires addressing the issue of color blindness. Furthermore, we found informally in experimenting with prototypes that a multiplicity of colors in the interface was often distracting and detracted from our ability to perform perceptual grouping for a given attribute. This anecdotal evidence, combined with the literature on the effects of mis-highlighting, convinced us to design for dynamic, single-color highlighting rather than for static, multi-color highlighting.

Keyword-in-Context Indices

Another alternative to our hi-cite design originates in the pioneering work of H.P. Luhn on keyword-in-context indices (KWIC indices) [4]. A typical keyword-in-context index has one entry for each keyword of each indexed work's title. These titles are presented "in context," meaning that they are presented sandwiched in between the surrounding words from the title. The keyword entries themselves are aligned in a single column. In a classic keyword-in-context index, there was a limit on how many characters could be displayed. Therefore, a fixed number of characters from the prefix to the keyword could be included to its left and a fixed number of characters from the keyword to the suffix could be included to its right. An example of a keyword-in-context index appears in Figure 4.

In KWIC indices, position is used to highlight a piece of structured information. "The initial letters of the alphabetized keywords form a column which guides the eye when scanning for desired words" [4, p. 289]. In other words, the "highlighting" mechanism of keyword-in-context is alignment rather than color. This reliance on alignment means that tables are more similar to keyword-in-context indices than to hi-cites.

In a static display, highlighting by alignment is easier for scanning than highlighting by color since the eye does not need to rove across the page. However, dynamically changing what is highlighted is much more disruptive for highlighting by alignment than for highlighting by color since it causes the rearrangement of all the visible document descriptions.

In the heterogeneous digital-library domain, users are likely to change what is highlighted quite often. It is for this reason that we hypothesize that hi-cites are better suited for skimming document descriptions in this environment than are KWIC indices. Hi-cites keep document descriptions globally constant in space throughout the browsing process while still allowing perceptual grouping by color.

```
ve Survey of the Radio, Phonograph, and Television Habits in
te Manual of the Edison Phonograph
                    The Phonograph and its Inventor, Thomas
ete Library of American Phonograph Recordings
    The Edison Improved Phonograph
           The Fabulous Phonograph; from Edison to Stereo
e Patent History of the Phonograph, 1877-1912: A Source Book
                        Recollections of Thomas A. Edison: A
```

Figure 4: A Keyword-in-Context Bibliographic Index

EVALUATION

Hi-cites are modeled after both citation sets and tables. Citation sets are easier to skim than tables since they take up less screen space and provide a more spatially continuous flow of information. Hi-cites are so close in appearance to citation sets that it is reasonable to believe that they are also easier to skim than tables. Tables are better than citation sets for the comparison of attribute values since they facilitate perceptual grouping. Hi-cites are different enough from tables that an evaluation is necessary to determine if they are also better than citation sets for this task. Before conducting such an evaluation, we enumerated the following hypotheses about the differences we expected to find among tables, citation sets, and hi-cites for tasks that involve looking for particular attributes within a set of document descriptions or for performing single attribute comparisons of document descriptions.

- Task-completion times will be fastest for tables, then hi-cites, then citation sets
- Perceived task-completion times will be fastest for tables, then hi-cites, then citation sets
- Error rates will be smallest for tables, then hi-cites, then citation sets
- Hi-cites will be the preferred condition since they allow for rapid answering of the questions as well as for quick skimming.

With these hypotheses in place, we constructed an experimental design that allowed us to determine their validity.

Experimental Design

Fourteen subjects participated in our evaluation of hi-cites. All subjects were from Stanford University and were paid for their time. Thirteen of the subjects were either undergraduates or graduate students. One subject was a Stanford staff member in the Computer Science Department. No subjects were students in the Computer Science Department, although the majority were students in the sciences.

The study compared three different presentation styles: tables, citation sets, and hi-cites. Each subject completed a practice task (where textual and oral help were provided), and then proceeded to complete two actual tasks (no help provided). Each task had three steps, where each step involved viewing a new collection of document descriptions in a particular presentation style and then answering questions about those descriptions. Step order and the pairings between collections and presentation styles were varied randomly. The questions asked in each step were of the following types (and were always presented in this fixed order):

- How many descriptions have the word X in the title?
- How many descriptions have publication locations listed?
- How many descriptions have the publication date Y?

All of these questions require that a user consider the document collection from the perspective of a particular attribute. We opted for providing users with this standard set of questions rather than providing a real-world task because we felt that these questions ensured that users would indeed perform attribute-based comparisons and obtain attribute-based overviews of the document collection (the focus of our evaluation, as well as a useful strategy in real digital-library tasks). Also, these question-based tasks were amenable to statistical comparison.

The time to complete each step and the error rate for each step were measured automatically. At the end of each task, subjects were also asked to rank each presentation style in terms of perceived speed. Finally, at the completion of the study, subjects were asked to compare/contrast orally the different styles and to choose one as preferred.

Experimental Results

Completion times

The comparison of completion times for each condition used the average of each subject's Task 1 time and Task 2 time. This average compensates for potential differences in the ease of learning each step, since step order was varied randomly for each subject. A one-way repeated-measures ANOVA on these average completion times ($F = 7.52$, $p = .003$) showed that the three conditions differed significantly at the .005 significance level. At significance level .05, the *post hoc* Student-Newman-Keuls test showed that the average completion times for tables and hi-cites were not significantly different. However, the average completion times for tables and citation sets and for hi-cites and citation sets were significantly different. Thus, tables and hi-cites are significantly faster for these tasks than are citation sets. The completion times for each condition, averaged across all subjects, are presented in Table 1.

Table 1: Average Completion Times (in seconds)

Tables	Citation Sets	Hi-Cites
43.081	51.969	45.477

Error rates

A one-way repeated-measures ANOVA on the error rates (again, averaged over the two tasks) showed no significant difference among the error rates ($F = .24$, $p = .786$). The error rates for each condition, averaged across all subjects, are presented in Table 2.

Table 2: Average Error Rates (number of wrong answers ÷ 3)

Tables	Citation Sets	Hi-Cites
0.10	0.13	0.11

Rankings of perceived completion times

In addition to looking at the actual differences in average completion times for each of the conditions, we also looked

at perceived differences in completion times. Users' perceptions of time do not always match up to real time. Furthermore, a user's perception that an interface is fast and facile can make a crucial difference in whether or not an interface will be used, no matter what the actual performance metrics might be for the interface.

A one-way repeated-measures ANOVA* ($F = 10.72$, $p < .005$) showed that the means for the average rankings (1 is fastest; 3 is slowest) were not equal. At a significance level of .01, the Student-Newman-Keuls test showed that all differences were significant. Hi-cites were perceived to be faster than tables, and tables were perceived to be faster than citation sets. The average rankings for each condition are shown in Table 3.

Table 3: Average Rankings of Perceived Completion Times (1 is fastest; 3 is slowest)

Tables	Citation Sets	Hi-Cites
1.82	2.68	1.50

Preferences

Subjects specified their overall preferred presentation style at the end of the session, not at the end of each task. This data was analyzed using a t-test, where the null hypothesis was that the probability of a subject choosing a condition as preferred was 1/3. The observed proportions of preference assignments for each condition appear in Table 4. Computing t for each condition showed that the observed proportions of preference assignments for hi-cites and citation sets were significantly different from 1/3 at the .05 level, but that the observed proportion for tables was not significantly different from 1/3 at the .05 level.

Table 4: Observed Proportions of Preference Assignments

Tables	Citation Sets	Hi-Cites
4/14	0/14	10/14

User feedback

At the end of a session, each subject was asked not only to select a presentation style as preferred, but also to compare and contrast all of the presentation styles. This feedback served to contextualize the preference choices with a concrete sense of why subjects found their preferred style to be useful and also gave a good indication of what was easy or hard to do in the other presentation styles. Overwhelmingly, subjects were enthusiastic about the value of hi-cites. They particularly liked the way color guided their eye to the "right place" in the collection. However, they found much to like about tables also. They pointed out that tables arranged the information in a format that would be very convenient to browse if it all fit on the screen. The need to scroll horizontally in tables was the most frequent complaint about their use. Subjects explained their difficulties with citation sets in terms of problems with honing in on the right piece of information in a description. They felt that color and space served to mark attributes well, but that slight changes in format and relative position were not good perceptual markers. A sampling of subject comments appears in Table 5.

Table 5: Sampling of Subject Comments

Hi-cites	caught my eye; much easier to go by color than location
	really great; focus on particular area; red drew visual attention
	with a lot of text, it just hops out at you
	stuck out compared to other ones
	don't know where to find first one (or there might not be a location)
Tables	hard to use because you need to scroll
	liked having it all mapped out
	nice; knew where to look
	maybe easier except you can't see it all on the screen
	everything structured in such a way you could do the same thing easily
Citation Sets	confusing; hard to pick out information being asked for
	more time to sort out information
	didn't really like them; just slower to get through
	tedious to read through words to look for information
	had to look through everything

* Technically, the use of the ANOVA is problematic here since the observations are not independent. However, the ANOVA is still the most useful statistic for this situation to the best of our knowledge.

Discussion

Before conducting the experiment, we set forth the following hypotheses about tables, citation sets, and hi-cites.

- Task-completion times will be fastest for tables, then hi-cites, then citation sets
- Perceived task-completion times will be fastest for tables, then hi-cites, then citation sets
- Error rates will be smallest for tables, then hi-cites, then citation sets
- Hi-cites will be the preferred condition since they allow for rapid answering of the questions as well as for quick skimming.

Statistical analysis of our experimental results leads to the following conclusions:.

- Task-completion times are faster for tables and hi-cites than for citation sets, but there is no significant difference between tables and hi-cites
- Perceived task-completion times are faster for hi-cites, then tables, then citation sets
- Error rates do not differ significantly among tables, citation sets, and hi-cites
- Hi-cites are the preferred condition

Our experimental results were somewhat surprising in that hi-cites either outranked tables or were not sufficiently different from tables in all cases. Thus, we can conclude that hi-cites are indeed a hybrid between tables and citation sets in that they:

- allow for ease of comparison
- allow for ease of skimming

Ease of comparison is demonstrated by our experimental results. Ease of skimming stems from the fact that hi-cites and citation sets are both compact representations of information that take up less screen space and are more spatially continuous than tables.

CONCLUSION

We have proposed the concept of hi-cites: dynamically created citations with active highlighting. The generation of hi-cites relies on access to attribute-value encodings of documents. The appearance of hi-cites is governed by citation styles that dictate what attributes should be presented, in what order they should be listed, and how they should be rendered and marked. Users control the active highlighting of hi-cites through a “tool tip” affordance. Point-plus-delay on any attribute value causes all values for that attribute to be highlighted in the citation set. We have chosen red color highlighting for hi-cites based upon our understanding of previous highlighting studies.

We have argued that hi-cites combine the benefits of both bibliographic citations and tables. The evaluation that we have performed supports our claim for a particular set of tasks and conditions. More insight into hi-cites could be gained by devising experiments to compare hi-cites to KWIC indices, to measure the ease of discovery for the point-plus-delay affordance currently used for hi-cites, and to determine the effects of different styles of highlighting (e.g., reverse video vs. red) in this context.

This paper has focused on the role of hi-cites in SenseMaker, an interface to a heterogeneous digital library. We believe that hi-cites can be useful in other environments that require heterogeneous structured descriptions to be browsed and compared with ease. Examples beyond digital libraries include product catalogs, classified advertisements, and World Wide Web search engines.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under Cooperative Agreement IRI-9411306. Funding for this agreement is also provided by DARPA, NASA, and the industrial partners of the Stanford Digital Libraries Project.

We would like to thank Cliff Nass and Vicky Reich for their valuable advice on experimental design and statistical analysis.

REFERENCES

1. Baldonado, M., Chang, C. K., Gravano, L., and Paepcke, A. Metadata for digital libraries: architecture and design rationale, in *Proceedings of DL ‘97* (Philadelphia PA, July 1997) ACM Press, 47-56.

2. Baldonado, M.Q.W., and Winograd, T. SenseMaker: An information-exploration interface supporting the contextual evolution of a user’s interests, in *Proceedings of CHI ‘97 (*Atlanta GA, March 1997) ACM Press, 11-18.

3. Brown, T.J. Visual display highlighting and information extraction, in *Proceedings of the Human Factors Society 35th Annual Meeting* (1991), 1427-1431.

4. Luhn, H.P. Keyword-in-context index for technical literature (KWIC Index). *American Documentation* XI, 4 (1960), 288-295.

5. Nielsen, J. *Usability Engineering*. Academic Press, Boston, MA, 1993.

6. Philipsen, G. Effects of six different highlighting modes on visual search performance in menu options. *International Journal of Human-Computer Interaction* 6, 3 (1994), 319-334.

7. Rheinfrank, J., and Evenson, S. Design languages, in T. Winograd (ed.), *Bringing Design to Software*. Addison-Wesley, Reading, MA, 1996.

8. Sellen, A., and Nicol, A. Building user-centered online help, in B. Laurel (ed.), *The Art of Human-Computer Interface Design*. Addison-Wesley, Reading, MA, 1990.

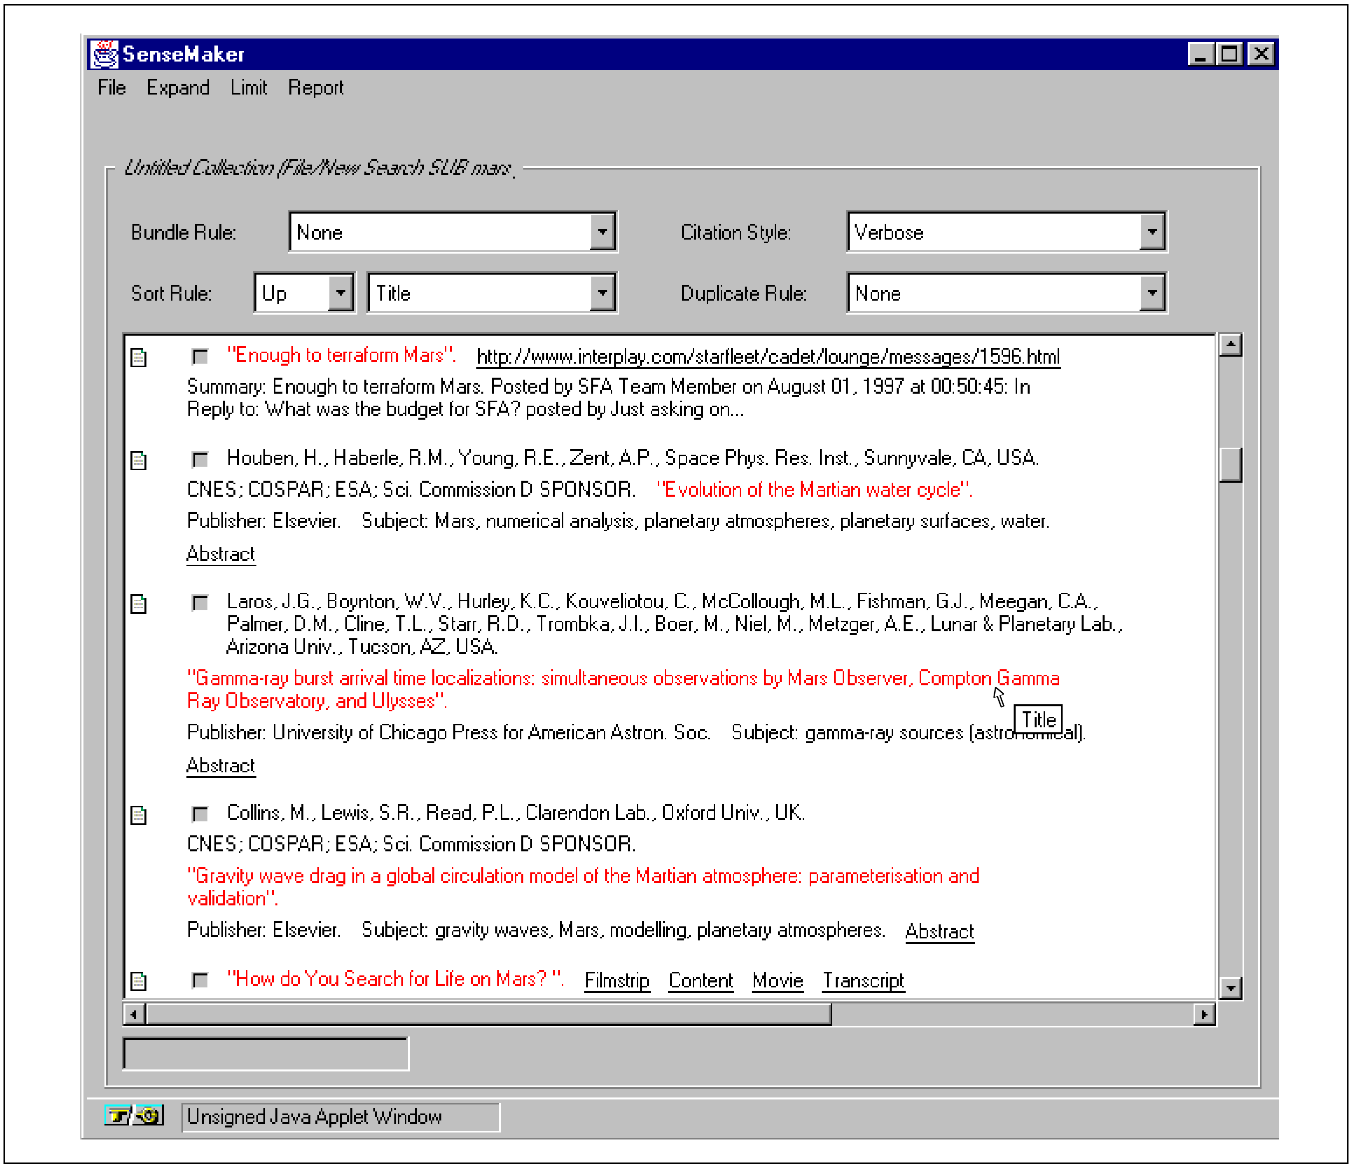

Figure 5: Hi-Cite Version of SenseMaker (*titles* are highlighted in red)
The heterogeneous set of documents described includes articles, WWW pages, and videos.

Hi-Cites:
Dynamically Created Citations with Active Highlighting

Michelle Q Wang Baldonado*
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304 USA
baldonad@parc.xerox.com

Terry Winograd
Gates 3B
Stanford University
Stanford, CA 94305-9035 USA
winograd@cs.stanford.edu

ABSTRACT

The original SenseMaker interface for information exploration [2] used tables to present heterogeneous document descriptions. In contrast, printed bibliographies and World Wide Web (WWW) search engines use formatted citations to convey this information. In this paper, we discuss *hi-cites*, a new interface construct developed for SenseMaker that combines the benefits of tables (which encourage the comparison of descriptions) and citations (which facilitate browsing).

Hi-cites are dynamically created citations with active highlighting. They are useful in environments where heterogeneous structured descriptions must be browsed and compared with ease. Examples beyond digital libraries include product catalogs, classified advertisements, and WWW search engines.

We have performed an evaluation of hi-cites, tables, and citations for tasks involving single attribute comparisons in the digital-library domain. This evaluation supports our claim that hi-cites are valuable for both comparison and skimming tasks in this environment.

Keywords

Hi-cites, dynamic citations, highlighting, digital libraries, information visualization, browsing

INTRODUCTION

Look for a book entitled "Physics" in most library card catalogs and you will be amazed at the number of books bearing that name. Fortunately, library catalogers describe documents by recording values for a wide variety of characteristic attributes. Library patrons can learn the author, publication year, number of pages, and the Library of Congress subject heading for each of the "Physics" books that have entries in a library's card catalog.

Users who perform information exploration in digital libraries spend much time perusing document descriptions of the kind found in library card catalogs. However, the descriptions come from an ever-growing variety of sources. Information may come not only from on-line library catalogs but also from video archives, World Wide Web (WWW) search engines, music indices, e-mail archives, map databases, and much more. The structured descriptions obtained from these sources can be heterogeneous in two ways. First, these sources may employ different conventions for describing documents. For example, USMARC defines one set of descriptive attributes, while Z39.50 Bib-1 defines a different set. In the Stanford Digital Library Project, we have addressed this problem through the development of a metadata architecture [1] that accommodates different descriptive attribute sets as well as services for translating descriptions from one attribute set to another.

A second type of document-description heterogeneity revolves around document genre. The choice of attributes for a document's description often depends on its genre. Document descriptions for music recordings are likely to include information about performers, tempo, and style; descriptions for WWW pages are likely to include information about URL and file size. The SenseMaker interface has been designed with this second style of heterogeneity in mind. The overall goal of SenseMaker is to support the contextual evolution of a user's interests during information exploration in the digital library (see [2] for more details). This paper focuses solely on how to present the heterogeneous document descriptions obtained while interacting with SenseMaker.

TABLES AND CITATION SETS

The original SenseMaker interface presented search results (document descriptions) in a table whose columns were dynamically defined by users. Tables encourage users to compare descriptions in terms of a particular attribute, due to our visual disposition to use whitespace and proximity as perceptual grouping cues [5]. Focusing on a single attribute while browsing a collection allows users to gain an overview of the collection with respect to that attribute. For example, focusing on *publication location* while browsing a collection of results produced by a search for "ODA" might lead the user to observe that much work on ODA has taken place in Europe. This style of overview is particularly useful in heterogeneous environments where document descriptions come from widely varying contexts. In a pilot user study performed on the original SenseMaker interface (in which users gathered references for a hypothetical term

Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '97), Los Angeles, California, April 1998.

* When this paper was first submitted, the author was a Ph.D. student at Stanford University.

paper on a given topic), users explicitly mentioned this ability to see all values for a particular attribute at once as one of the features that they particularly liked.

Figure 1 shows an excerpt from the original table-based version of SenseMaker. Each row in the table corresponds to one document. Each column in the table stands for an individual attribute of the document.

Figure 1: Excerpt from Table-Based Version of SenseMaker

As this figure illustrates, a tabular view of heterogeneous document descriptions is often sparse. We have *URL* values but no *author* values for results from WWW search engines, and vice versa from citation databases. The multiplicity of "*Not available*" values is distracting to users. Furthermore, users must almost always use horizontal scrolling since screen space disappears quickly when each attribute requires its own column. Skimming is also difficult due to the spatial discontinuities that are characteristic of tables.

In contrast to tables, formatted bibliographies require less screen space and provide a spatially continuous flow of information. A bibliographic citation concatenates just the attribute values that are available for the document it describes. Furthermore, these citations make use of conventions that indicate for each value what its associated attribute is. These conventions may rely on fixed order (most bibliographic styles prescribe a particular order for citation elements) and may also use visual cues such as font style. A typical citation appears in Figure 2.

> Strunk, W. Jr. and White, E. B. *The Elements of Style*. Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Figure 2: Typical Citation

In this figure, italics are used for book titles, while the order of author, book title, edition, publisher, publication city, and publication year is conventional. The demarcating punctuation marks between items serve to delineate and identify each item in the citation as well.

In the publishing world, standard bibliographic styles are used for traditional written works, such as books, magazine articles, newspapers, memos and more. Formatted bibliographies can be generalized to accommodate nontraditional document descriptions as well (e.g., descriptions of videos, software, e-mail, maps, people, etc.). We call such a generalized bibliography a *citation set*. Note that most World Wide Web search engines present their results as citation sets (see Figure 3) .

Figure 3: AltaVista Citations

The gain in browsing ease offered by citation sets (due to the compactness and visually identifiable structure of citations) is offset by a loss in ease of comparison of available items. In this paper, we define and propose *hi-cites*, a novel interface idea that promises to improve browsing by combining the benefits of citation sets with the benefits of tables. The SenseMaker interface has been revised to make use of hi-cites instead of tables.

HI-CITES

Hi-cites, dynamically created citations with active highlighting, are a hybrid of citation sets and tables. They are like citation sets in that they concatenate visually marked attribute values (subject to wraparound), treating attribute values as though they were words and phrases in a single sentence. Definable citation styles encapsulate the ordering, marking, and rendering rules required for this display. For example, one citation style might stipulate that *title* appear before *author*, while another might change that ordering. In addition, one citation style might mark an *editor* value with the suffix "ed.," while another might use the suffix "(editor)." Finally, a citation style might specify that article *titles* are quoted while journal *titles* are italicized.

Hi-cites are like tables in that they facilitate the perceptual grouping of all values of a particular attribute. Specifically, users can highlight* in red all values for an attribute by pausing for a short period of time over that attribute with the mouse (see the next section for a detailed discussion of this affordance). Figure 4 shows both *title* highlighting and *publisher* highlighting. Note that it substitutes reverse video for color so that the highlighting can be represented in a black-and-white figure.

In addition to enabling highlighting, the point-plus-delay affordance also allows users to find out the name of each attribute. If a SenseMaker user were to move her cursor over a *title* attribute value and pause, then all *title* values would be highlighted. In addition, the name *Title* would

* The name hi-cite derives from the words "highlight" and "citation."

Strunk, W. Jr. and White, E. B. *The Elements of Style*. Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Morris, W. and Morris, M. *Morris Dictionary of Word and Phrase Origins*. Harper and Row, Publishers, New York, 1977.

Strunk, W. Jr. and White, E. B. *The Elements of Style*. Third edition. Macmillan Publishing Co., Inc., New York, 1979.

Morris, W. and Morris, M. *Morris Dictionary of Word and Phrase Origins*. Harper and Row, Publishers, New York, 1977.

Figure 4: Two Sets of Hi-Cites (actual hi-cites make use of color rather than reverse video for highlighting)

show up in a temporary rectangular pop-up box. This result is shown in Figure 5.

This "help" feature is important for cases where users are not familiar with the chosen attribute marking conventions. For example, a newcomer to geographic information sources may not be familiar with the conventional marking of latitude and longitude, and thus may need to discover what the pieces of a geographic citation are.

Affordances

The hi-cite affordance for highlighting an attribute and for discovering an attribute's name (point-plus-delay over an individual attribute value) is similar to that needed to initiate "tool tips" in Windows applications. The choice of this affordance was made after comparing it to a number of other possibilities. In an early design that we considered, the affordance was simply moving the mouse over an attribute value, similar to the affordance for Macintosh "balloon help" [8]. However, we discovered that this was distracting to the user. Furthermore, the interaction between scrolling and skimming was problematic. Consider the situation where a user wants to skim all titles in the collection of document descriptions (not just those that fit into the current window). If moving the mouse over an attribute value causes highlighting to occur, then the act of moving the mouse to the scrollbar is likely to cause the desired highlighting to be lost.

We rejected the possibilities of highlighting on single-click or double-click on the grounds that they would conflict with the design languages [7] of Web browsers and GUIs, both of which are already well-known and familiar to our intended user community. In particular, single-click maps onto "follow link" in Web browsers, while double-click maps onto "open" in GUIs.

Finally, we debated making the affordance something less familiar, such as typing a letter over the attribute value (e.g., "h" for highlight); right-clicking and then either automatically highlighting or bringing up a pop-up menu (as in Windows); or doing a combination of special key (e.g., the Apple key or the Control key) and clicking. In all of these cases, the affordance would be learnable, but it was unlikely that users would "stumble" across the affordance. We felt that highlighting was so crucial to hi-cites that ease of discovery was as important of ease of learning. Thus, we decided to use the point-plus-delay affordance, even though the delay means that highlighting is not as fast as it could be for the user. Clearly, choice of delay time deserves further study.

Implementing Citation Styles

Many different citation styles exist in the publishing world. Hi-cites also allow for different citation styles. In SenseMaker, a citation style is implemented as an ordered list of attributes. In addition, objects specify how each attribute should be rendered (e.g., italicized or quoted) and marked (e.g., with a suffix "ed."). In our implementation, attribute values can be rendered as links (pointers) to the actual attribute values. This use of links in citations has already become common practice in the reference sections of documents on the World Wide Web. For example, articles in the on-line D-Lib magazine (http://www.dlib.org/) often include bibliographic citations that have hypertext links to the referenced works. Similarly, the result format for the various WWW search engines is an implicit citation style that contains links.

Including link information in the rendering-style object also allows local values that are particularly long to be accessible via a link rather than directly integrated into the citation. Consider for example the difference in expected lengths of a title, abstract, and full text of a document. Titles are usually fairly short in length and are almost always listed in on-line citations. Abstracts are usually on the order of one to two paragraphs. In many electronic card catalogs, the user has a choice of a terse or verbose presentation of citations, where the abstract is listed in only the verbose style. Finally, full text is almost never included directly in the citations of an on-line catalog, but may be available via a special command. The ability to have links in citations changes the design space in that it makes it possible to have longer values readily accessible without giving up screen space.

In non-textual domains, such as video, the "citation" for a video clip may be a multimedia sequence of images and text rather than an attribute-value style citation. In general, the question of how to display succinct representations of multimedia documents is an open area for research. We believe that at least some of the hi-cite principles are applicable to this domain. For example, the last hi-cite that appears in Figure 5 is a description of a video. It contains a *title* attribute, which is common to all of the displayed hi-cites, but it also contains a number of genre-specific attributes (e.g., *filmstrip*) that are accessible via links. An alternative to this citation style might inline the *filmstrip* value for each video directly into the hi-cite.

RELATED WORK

Highlighting

The use of highlighting as a mechanism for bringing a user's attention to a particular piece of information on the screen has been carefully studied by the human factors community. From this perspective, highlighting is defined as any visual characteristic that causes the highlighted piece of In addition to color, highlighting can include boxing, inverse video, blinking, and other techniques. Several studies have shown that color highlighting is useful for the visual-search task [3, 6]. Although no definitive answer has emerged on how to choose a color for highlighting, some studies show the perceptual salience of red [6]. In addition to choosing a color for highlighting, designers must also consider the amount of color in the interface.

One issue that has received great attention in highlighting studies is the effect of mis-highlighting information, or what happens when information is highlighted but the user does not wish to focus on it. In general, a high rate of mis-highlighting can override the benefits obtained from highlighting [6]. These findings are good evidence for the importance of user control in hi-cites. The fact that a hi-cite user can quickly and easily control what piece of information is highlighted means that the user can ensure that highlighting always matches the current focus of attention. These study results also indicate the need for SenseMaker users to be able to turn off highlighting completely in the interface, a feature that is not currently implemented.

One possible alternative to hi-cites is to differentiate the attributes within a citation by displaying each of them in a different color (e.g., red for *title*, blue for *publication date*). In this case, dynamic highlighting would not be necessary. Users could compare values for an attribute by focusing on the color corresponding to that attribute. However, this strategy requires addressing the issue of color blindness. Furthermore, we found informally in experimenting with prototypes that a multiplicity of colors in the interface was often distracting and detracted from our ability to perform perceptual grouping for a given attribute. This anecdotal evidence, combined with the literature on the effects of mis-highlighting, convinced us to design for dynamic, single-color highlighting rather than for static, multi-color highlighting.

Keyword-in-Context Indices

Another alternative to our hi-cite design originates in the pioneering work of H.P. Luhn on keyword-in-context indices (KWIC indices) [4]. A typical keyword-in-context index has one entry for each keyword of each indexed work's title. These titles are presented "in context," meaning that they are presented sandwiched in between the surrounding words from the title. The keyword entries themselves are aligned in a single column. In a classic keyword-in-context index, there was a limit on how many characters could be displayed. Therefore, a fixed number of characters from the prefix to the keyword could be included to its left and a fixed number of characters from the keyword to the suffix could be included to its right. An example of a keyword-in-context index appears in Figure 4.

In KWIC indices, position is used to highlight a piece of structured information. "The initial letters of the alphabetized keywords form a column which guides the eye when scanning for desired words" [4, p. 289]. In other words, the "highlighting" mechanism of keyword-in-context is alignment rather than color. This reliance on alignment means that tables are more similar to keyword-in-context indices than to hi-cites.

In a static display, highlighting by alignment is easier for scanning than highlighting by color since the eye does not need to rove across the page. However, dynamically changing what is highlighted is much more disruptive for highlighting by alignment than for highlighting by color since it causes the rearrangement of all the visible document descriptions.

In the heterogeneous digital-library domain, users are likely to change what is highlighted quite often. It is for this reason that we hypothesize that hi-cites are better suited for skimming document descriptions in this environment than are KWIC indices. Hi-cites keep document descriptions globally constant in space throughout the browsing process while still allowing perceptual grouping by color.

```
ve Survey of the Radio, Phonograph, and Television Habits in
te Manual of the Edison Phonograph
                    The Phonograph and its Inventor, Thomas
ete Library of American Phonograph Recordings
    The Edison Improved Phonograph
           The Fabulous Phonograph; from Edison to Stereo
e Patent History of the Phonograph, 1877-1912: A Source Book
                        Recollections of Thomas A. Edison: A
```

Figure 4: A Keyword-in-Context Bibliographic Index

EVALUATION

Hi-cites are modeled after both citation sets and tables. Citation sets are easier to skim than tables since they take up less screen space and provide a more spatially continuous flow of information. Hi-cites are so close in appearance to citation sets that it is reasonable to believe that they are also easier to skim than tables. Tables are better than citation sets for the comparison of attribute values since they facilitate perceptual grouping. Hi-cites are different enough from tables that an evaluation is necessary to determine if they are also better than citation sets for this task. Before conducting such an evaluation, we enumerated the following hypotheses about the differences we expected to find among tables, citation sets, and hi-cites for tasks that involve looking for particular attributes within a set of document descriptions or for performing single attribute comparisons of document descriptions.

- Task-completion times will be fastest for tables, then hi-cites, then citation sets
- Perceived task-completion times will be fastest for tables, then hi-cites, then citation sets
- Error rates will be smallest for tables, then hi-cites, then citation sets
- Hi-cites will be the preferred condition since they allow for rapid answering of the questions as well as for quick skimming.

With these hypotheses in place, we constructed an experimental design that allowed us to determine their validity.

Experimental Design

Fourteen subjects participated in our evaluation of hi-cites. All subjects were from Stanford University and were paid for their time. Thirteen of the subjects were either undergraduates or graduate students. One subject was a Stanford staff member in the Computer Science Department. No subjects were students in the Computer Science Department, although the majority were students in the sciences.

The study compared three different presentation styles: tables, citation sets, and hi-cites. Each subject completed a practice task (where textual and oral help were provided), and then proceeded to complete two actual tasks (no help provided). Each task had three steps, where each step involved viewing a new collection of document descriptions in a particular presentation style and then answering questions about those descriptions. Step order and the pairings between collections and presentation styles were varied randomly. The questions asked in each step were of the following types (and were always presented in this fixed order):

- How many descriptions have the word X in the title?
- How many descriptions have publication locations listed?
- How many descriptions have the publication date Y?

All of these questions require that a user consider the document collection from the perspective of a particular attribute. We opted for providing users with this standard set of questions rather than providing a real-world task because we felt that these questions ensured that users would indeed perform attribute-based comparisons and obtain attribute-based overviews of the document collection (the focus of our evaluation, as well as a useful strategy in real digital-library tasks). Also, these question-based tasks were amenable to statistical comparison.

The time to complete each step and the error rate for each step were measured automatically. At the end of each task, subjects were also asked to rank each presentation style in terms of perceived speed. Finally, at the completion of the study, subjects were asked to compare/contrast orally the different styles and to choose one as preferred.

Experimental Results

Completion times

The comparison of completion times for each condition used the average of each subject's Task 1 time and Task 2 time. This average compensates for potential differences in the ease of learning each step, since step order was varied randomly for each subject. A one-way repeated-measures ANOVA on these average completion times ($F = 7.52$, $p = .003$) showed that the three conditions differed significantly at the .005 significance level. At significance level .05, the *post hoc* Student-Newman-Keuls test showed that the average completion times for tables and hi-cites were not significantly different. However, the average completion times for tables and citation sets and for hi-cites and citation sets were significantly different. Thus, tables and hi-cites are significantly faster for these tasks than are citation sets. The completion times for each condition, averaged across all subjects, are presented in Table 1.

Table 1: Average Completion Times (in seconds)

Tables	Citation Sets	Hi-Cites
43.081	51.969	45.477

Error rates

A one-way repeated-measures ANOVA on the error rates (again, averaged over the two tasks) showed no significant difference among the error rates ($F = .24$, $p = .786$). The error rates for each condition, averaged across all subjects, are presented in Table 2.

Table 2: Average Error Rates (number of wrong answers ÷ 3)

Tables	Citation Sets	Hi-Cites
0.10	0.13	0.11

Rankings of perceived completion times

In addition to looking at the actual differences in average completion times for each of the conditions, we also looked

at perceived differences in completion times. Users' perceptions of time do not always match up to real time. Furthermore, a user's perception that an interface is fast and facile can make a crucial difference in whether or not an interface will be used, no matter what the actual performance metrics might be for the interface.

A one-way repeated-measures ANOVA[*] (F = 10.72, p < .005) showed that the means for the average rankings (1 is fastest; 3 is slowest) were not equal. At a significance level of .01, the Student-Newman-Keuls test showed that all differences were significant. Hi-cites were perceived to be faster than tables, and tables were perceived to be faster than citation sets. The average rankings for each condition are shown in Table 3.

Table 3: Average Rankings of Perceived Completion Times (1 is fastest; 3 is slowest)

Tables	Citation Sets	Hi-Cites
1.82	2.68	1.50

Preferences

Subjects specified their overall preferred presentation style at the end of the session, not at the end of each task. This data was analyzed using a t-test, where the null hypothesis was that the probability of a subject choosing a condition as preferred was 1/3. The observed proportions of preference assignments for each condition appear in Table 4. Computing t for each condition showed that the observed proportions of preference assignments for hi-cites and citation sets were significantly different from 1/3 at the .05 level, but that the observed proportion for tables was not significantly different from 1/3 at the .05 level.

Table 4: Observed Proportions of Preference Assignments

Tables	Citation Sets	Hi-Cites
4/14	0/14	10/14

User feedback

At the end of a session, each subject was asked not only to select a presentation style as preferred, but also to compare and contrast all of the presentation styles. This feedback served to contextualize the preference choices with a concrete sense of why subjects found their preferred style to be useful and also gave a good indication of what was easy or hard to do in the other presentation styles. Overwhelmingly, subjects were enthusiastic about the value of hi-cites. They particularly liked the way color guided their eye to the "right place" in the collection. However, they found much to like about tables also. They pointed out that tables arranged the information in a format that would be very convenient to browse if it all fit on the screen. The need to scroll horizontally in tables was the most frequent complaint about their use. Subjects explained their difficulties with citation sets in terms of problems with honing in on the right piece of information in a description. They felt that color and space served to mark attributes well, but that slight changes in format and relative position were not good perceptual markers. A sampling of subject comments appears in Table 5.

Table 5: Sampling of Subject Comments

Hi-cites	caught my eye; much easier to go by color than location
	really great; focus on particular area; red drew visual attention
	with a lot of text, it just hops out at you
	stuck out compared to other ones
	don't know where to find first one (or there might not be a location)
Tables	hard to use because you need to scroll
	liked having it all mapped out
	nice; knew where to look
	maybe easier except you can't see it all on the screen
	everything structured in such a way you could do the same thing easily
Citation Sets	confusing; hard to pick out information being asked for
	more time to sort out information
	didn't really like them; just slower to get through
	tedious to read through words to look for information
	had to look through everything

* Technically, the use of the ANOVA is problematic here since the observations are not independent. However, the ANOVA is still the most useful statistic for this situation to the best of our knowledge.

Discussion

Before conducting the experiment, we set forth the following hypotheses about tables, citation sets, and hi-cites.

- Task-completion times will be fastest for tables, then hi-cites, then citation sets
- Perceived task-completion times will be fastest for tables, then hi-cites, then citation sets
- Error rates will be smallest for tables, then hi-cites, then citation sets
- Hi-cites will be the preferred condition since they allow for rapid answering of the questions as well as for quick skimming.

Statistical analysis of our experimental results leads to the following conclusions:.

- Task-completion times are faster for tables and hi-cites than for citation sets, but there is no significant difference between tables and hi-cites
- Perceived task-completion times are faster for hi-cites, then tables, then citation sets
- Error rates do not differ significantly among tables, citation sets, and hi-cites
- Hi-cites are the preferred condition

Our experimental results were somewhat surprising in that hi-cites either outranked tables or were not sufficiently different from tables in all cases. Thus, we can conclude that hi-cites are indeed a hybrid between tables and citation sets in that they:

- allow for ease of comparison
- allow for ease of skimming

Ease of comparison is demonstrated by our experimental results. Ease of skimming stems from the fact that hi-cites and citation sets are both compact representations of information that take up less screen space and are more spatially continuous than tables.

CONCLUSION

We have proposed the concept of hi-cites: dynamically created citations with active highlighting. The generation of hi-cites relies on access to attribute-value encodings of documents. The appearance of hi-cites is governed by citation styles that dictate what attributes should be presented, in what order they should be listed, and how they should be rendered and marked. Users control the active highlighting of hi-cites through a “tool tip” affordance. Point-plus-delay on any attribute value causes all values for that attribute to be highlighted in the citation set. We have chosen red color highlighting for hi-cites based upon our understanding of previous highlighting studies.

We have argued that hi-cites combine the benefits of both bibliographic citations and tables. The evaluation that we have performed supports our claim for a particular set of tasks and conditions. More insight into hi-cites could be gained by devising experiments to compare hi-cites to KWIC indices, to measure the ease of discovery for the point-plus-delay affordance currently used for hi-cites, and to determine the effects of different styles of highlighting (e.g., reverse video vs. red) in this context.

This paper has focused on the role of hi-cites in SenseMaker, an interface to a heterogeneous digital library. We believe that hi-cites can be useful in other environments that require heterogeneous structured descriptions to be browsed and compared with ease. Examples beyond digital libraries include product catalogs, classified advertisements, and World Wide Web search engines.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under Cooperative Agreement IRI-9411306. Funding for this agreement is also provided by DARPA, NASA, and the industrial partners of the Stanford Digital Libraries Project.

We would like to thank Cliff Nass and Vicky Reich for their valuable advice on experimental design and statistical analysis.

REFERENCES

1. Baldonado, M., Chang, C. K., Gravano, L., and Paepcke, A. Metadata for digital libraries: architecture and design rationale, in *Proceedings of DL ‘97* (Philadelphia PA, July 1997) ACM Press, 47-56.

2. Baldonado, M.Q.W., and Winograd, T. SenseMaker: An information-exploration interface supporting the contextual evolution of a user’s interests, in *Proceedings of CHI ‘97 (*Atlanta GA, March 1997) ACM Press, 11-18.

3. Brown, T.J. Visual display highlighting and information extraction, in *Proceedings of the Human Factors Society 35th Annual Meeting* (1991), 1427-1431.

4. Luhn, H.P. Keyword-in-context index for technical literature (KWIC Index). *American Documentation* XI, 4 (1960), 288-295.

5. Nielsen, J. *Usability Engineering*. Academic Press, Boston, MA, 1993.

6. Philipsen, G. Effects of six different highlighting modes on visual search performance in menu options. *International Journal of Human-Computer Interaction* 6, 3 (1994), 319-334.

7. Rheinfrank, J., and Evenson, S. Design languages, in T. Winograd (ed.), *Bringing Design to Software*. Addison-Wesley, Reading, MA, 1996.

8. Sellen, A., and Nicol, A. Building user-centered online help, in B. Laurel (ed.), *The Art of Human-Computer Interface Design*. Addison-Wesley, Reading, MA, 1990.

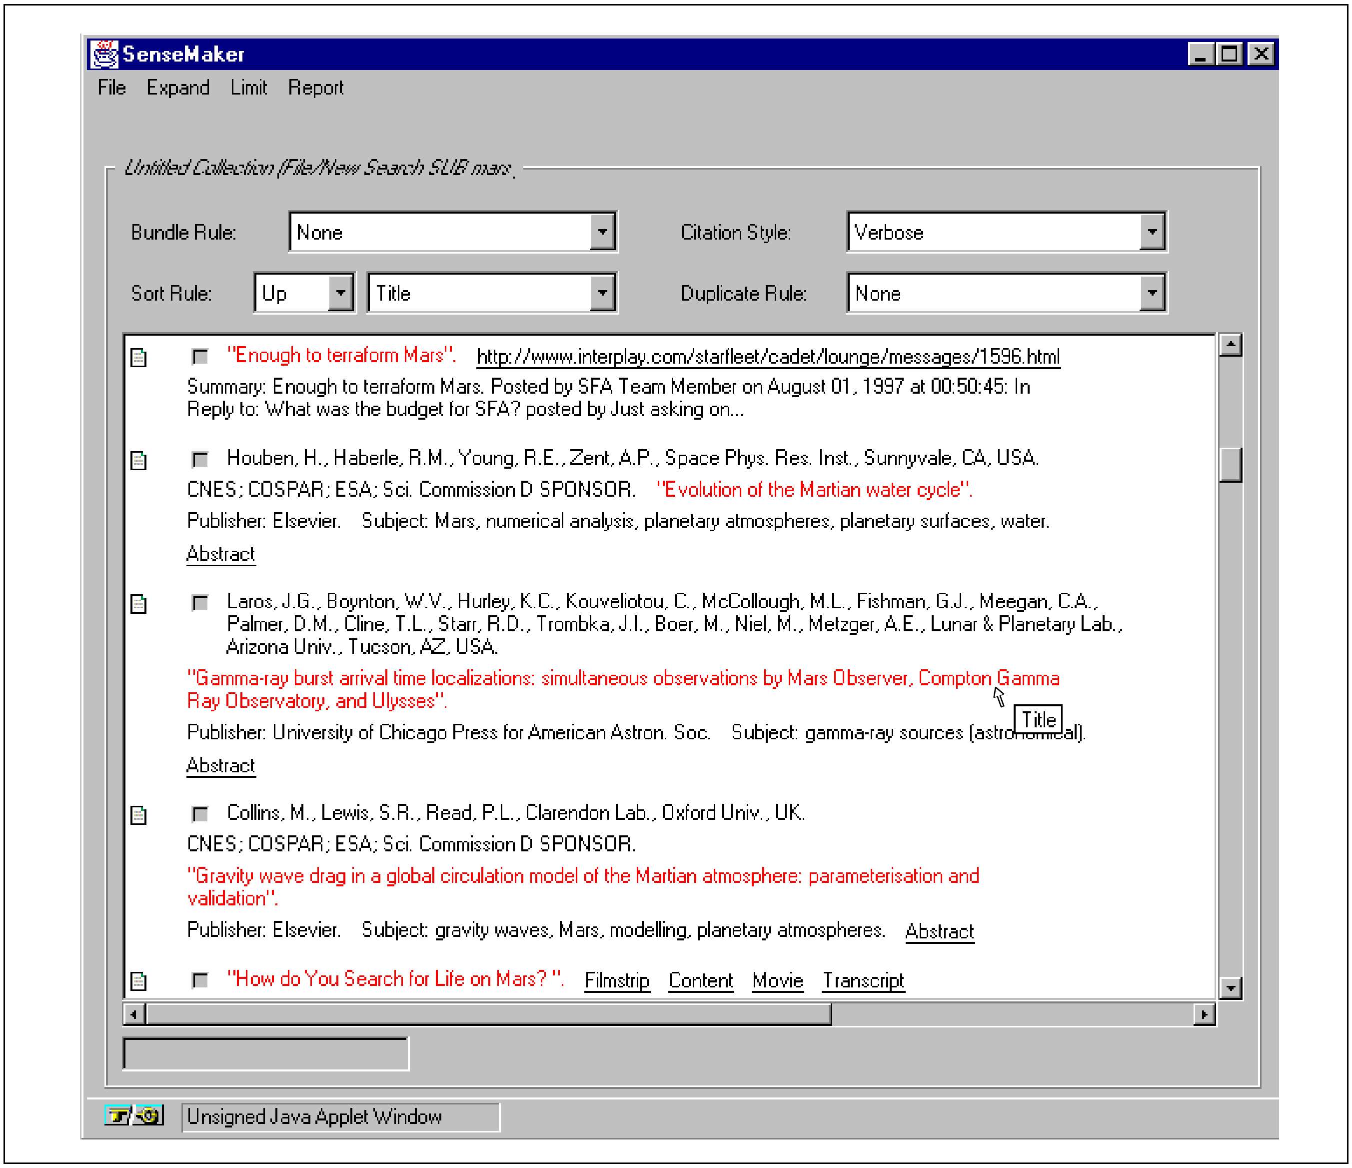

Figure 5: Hi-Cite Version of SenseMaker (*titles* are highlighted in red)
The heterogeneous set of documents described includes articles, WWW pages, and videos.