

Medical Sensemaking with Entity Workspace

Dorrit Billman and Eric A. Bier

Palo Alto Research Center, Inc.

3333 Coyote Hill Rd., Palo Alto, CA 94304

{dbillman, bier}@parc.com

ABSTRACT

Knowledge workers making sense of a topic divide their time among activities including searching for information, reading, and taking notes. We have built a software system that supports and integrates these activities. To test its effectiveness, we conducted a study where subjects used it to perform medical question-answering tasks. Initial results indicate that subjects could use the system, but that the nature of this use depended on the subject's overall question-answering strategy. Two dominant strategies emerged that we call the Reader and Searcher strategies.

Author Keywords

sensemaking, search, note-taking, quick click, highlighting.

ACM Classification Keywords

H.5.2 [Information interfaces and presentation (e.g., HCI)]: User Interfaces---GUI.

INTRODUCTION

Making sense of a rich body of information preparatory to acting on it requires integrating many activities: search for relevant information; selective reading; and capturing, commenting on, and organizing information in notes. Sensemaking is a central activity for many knowledge workers, including lawyers, doctors, scientists, researchers, intelligence analysts, detectives, financial analysts, industry analysts, and ordinary people making complex decisions.

Our software system, Entity Workspace [2], is designed to facilitate sensemaking by providing an integrated work environment. It supports search by providing a document-collection visualization and methods for searching over the collection and within documents. It supports reading with automatic highlighting of important terms. It supports note-taking with an electronic notebook and quick methods for importing text from documents, adding comments, and organizing information. For some tasks, this system takes

the place of multiple traditional tools, such as search engines, word processors, and document collection browsers. It has similarities to tools such as Oculus nSpace [4], but differs in its approach to these tasks. We assessed the effectiveness of our prototype in a small-scale study. Users worked on a medical sensemaking task in which they answered questions that varied in the amount of inference and integration across documents required of the user.

A TOOL TO ACCELERATE SENSEMAKING

We designed Entity Workspace (EW) to help users make sense of large document collections. Figure 1 shows EW's two, multi-pane windows as used in our study. While EW can be used on a single display, two were used to minimize window management and because multiple screens are becoming common in offices. The window on the left, known as CorpusView, provides an array visualization of a document collection, an area for reading documents (two documents are open in the Figure), and a region for specifying search queries that apply either across the entire document collection or within a single document.

The window on the right side of the figure provides a facility for taking structured notes. It has four panels. 1) In the largest, the Evidence Panel, users take notes about entities. 2) The Entity Inspector Panel displays information about a selected entity. 3) The System Suggestions Panel recommends additional entities of interest. 4) The Document Trails panel shows the history of searches and the documents found by each search.

Searching. EW provides tools to search the corpus for relevant documents, to search individual documents, and to track the process of using the documents identified by search. The design goal is to provide both efficiency and power. There are two methods to initiate search of the corpus: through a query box and as a side effect of entity selection. There is one method, a second query-box, for initiating search within the current active document. Both boxes support Boolean queries.

When a corpus level query is initiated, the system finds all matching documents and highlights them in a visualization that shows the collection as an array of icons. In addition, a modified version of this query seeds the document-level query box, so the user can immediately find instances of the search terms within each matching document. This method is designed for power.

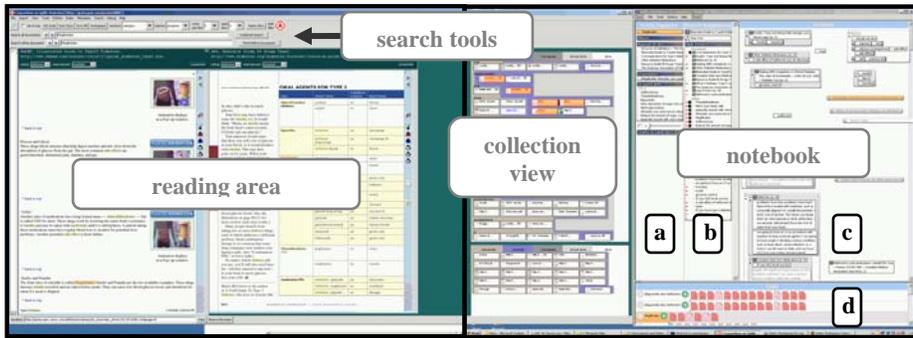


Figure 1. Experimental setup on three displays. Labels (a)-(d) indicate panels.

The second corpus-level search method, *one-click*, provides a low effort way to specify a search. When the user clicks on a notebook entity that represents a document page, EW opens the corresponding document to that page. If the user clicks on any other entity, EW computes a query to find that entity, enters that query in the corpus-level query box, runs the query, and highlights the found documents in the icon array. This method is designed for efficiency.

EW provides three methods for browsing within documents returned from a corpus-level search. First, arrow-shaped buttons beside each query box allow the user to view the matching documents in decreasing or increasing order of search score. Second, the user can also visit the matches in arbitrary order by clicking on them in the collection array. Finally, in the Document Trails panel, the tool displays a history of recent queries, showing icons that represent the matching documents. Clicking on an icon in a trail opens the matching document. Trails also provide arrow keys for stepping through the documents in order.

Reading. To accelerate reading, EW displays all documents as page images for rapid page flipping. It automatically highlights terms that the user has already placed in the notebook, using more salient highlighting for terms that the user has marked as being of high interest.

Note-taking. EW provides several methods for integrated importing, commenting, and organization of notes. The design goal is to capture more relevant information faster.

Importing. EW provides *entity quick-click* [1] to copy phrases or longer passages to the notebook with little effort. If the desired passage is already highlighted, it can be copied by depressing the Shift key and clicking. Otherwise, the user first selects the passage and then performs the Shift-click. In either case, the copied text is positioned in the notebook automatically; little mouse travel is needed. The copied phrase is automatically grouped with information about its source document and page.

Finally, to accelerate note organization, EW provides a drag-and-drop method for grouping phrases that we call *snap-together knowledge*. If one phrase is dragged to be near another horizontally, they snap together into a *belief statement*. When phrases or belief statements are positioned near each other vertically, they snap together

into an *entity group*. If two entity groups are positioned near each other vertically, they snap together into a *super group*. Clicking on an entity in the notebook causes all belief statements and entity groups that contain that entity to highlight, making it easy to review notes about that entity. Figure 2 shows an example super group including entity groups, which in turn contain document links, a belief statement, an entity, a comment and copied passages.

The notebook is designed to scale up to hundreds of notes. Groups and supergroups can be collapsed to a compact form and the notebook can be scrolled to fit more notes.

These technologies work together. Notebook entries generate document highlighting. Note selection initiates a search or opens a document, and creates a document trail. Organizing notes influences the information that will be highlighted when a note is selected and hence which entities and documents will be easy to find later.

MEDICAL SENSEMAKING USER STUDY

EW was originally designed for intelligence analysis. However, for this study we chose medical sensemaking tasks to verify that our technologies could be used in multiple domains and because such tasks are of general importance and interest. Initial study results indicate that novice subjects were able to take advantage of our system. While the interface is complex, users were able to learn the basic task functions in a half-hour training.

Further, the nature of their use depended on the subject's overall strategy. We noticed two primary strategy groups that we called Readers and Searchers. The study also provided us with guidance on ways to simplify our user interface and increase its performance and utility.

Method

Task. The medical task concerned diabetes. Three questions

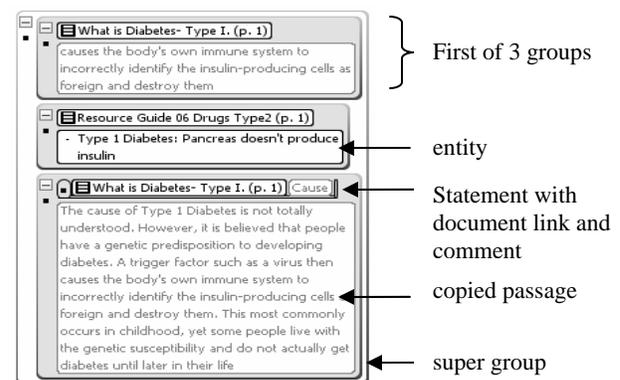


Figure 2. Excerpt from a subject's notebook

involved general information and five questions asked subjects to advocate for a hypothetical uncle newly diagnosed with diabetes. To motivate note-taking, subjects were asked to prepare information another family member might use when accompanying the uncle on a medical visit. Questions varied: (1) in how directly the information could be searched from key words in the question versus how much interpretation and integration was needed; (2) in whether information had to be collected across multiple documents or was available in one, and (3) in the relevant methods for assessing if an answer was complete. The questions were ordered so that information and documents found early might contribute to later question-answering.

Corpus. The 66 documents addressed various aspects of diabetes. Documents included web pages and a variety of papers. Formality ranged from papers written for clinicians or scientists to patient discussion postings on diabetes websites. Documents were classified into a dozen categories such as “intro”, “drugs”, and “exercise”. These categories were shown in the visualization.

Users. Nine employees, primarily student interns, volunteered for the study. None had used Entity Workspace before, although a few had heard it discussed.

Procedure. The study ran for two and a half to three hours. Users were encouraged to ask questions and talk about what they were doing. Sessions were videotaped and results were saved at intermediate points. Users answered a pre-assessment questionnaire concerning relevant health knowledge. Training used a different corpus with short documents concerning an intelligence analysis problem. It included an inference task as well as use of the tools and took slightly under an hour. Users then read the “uncle” scenario and began working with the diabetes corpus. The last six users spent roughly 5 minutes orienting to and exploring the corpus before they began answering questions. Users began with the 3 general questions from the preassessment questionnaire and continued with the five “uncle” questions. At the end of the session users were asked what they liked and didn’t like about the system.

Subjects were seated in front of a PC with three displays, all set at 1600 pixels wide by 1200 high. The application windows were set up as shown in Figure 1. In left to right order, subjects were given a region for reading documents, the document collection array, a notebook window, and a Web browser displaying a tutorial for our tools (which was used little, if at all).

The procedure was informal; the experimenter answered questions, made occasional suggestions (e.g., if the user was having trouble finding a function), and moved users through the tasks so we could observe subject responses to a range of tasks. Of nine users, three finished all questions.

Results

Overall, users were able understand the task and the tool, and to produce sensible, though often incomplete, answers

to the questions. The approach to the task and choice of methods from the tool both varied. We report on the strategies employed and then on tool usability.

User Strategies: Readers and Searchers

We noticed two approaches to the task. Subjects characterized as *Readers* spent more time reading the content of the articles. By and large, these users expected that they would need to understand or think about the material more and to integrate information from multiple documents. They were motivated to take notes on what they had learned so far, to keep track of previously-read documents so they could return to them later, and to add commentary. One Reader mentioned building and using a custom document corpus in the past.

Searchers, by contrast, expected to be able to find the answer to each task question by direct match to a search. With this strategy, users seemed to expect that they could quickly determine if the search had succeeded in finding (part of) the needed information. Hence, they did not expect to need to read or make inferences from the information. They did less reading and note-taking, focusing their energy on constructing the best possible search queries, in the hopes of finding hits with a clearly relevant and probably complete answer. Searchers may not have been familiar with the idea of having a pre-selected document corpus but were very familiar with using Web search engines. Searchers sometimes complained about not having access to the whole Web where, they believed, they could easily find the requested information. Two subjects adopted a strong Searcher strategy, three adopted a Reader strategy, with the remaining four in between.

While other authors have reported on alternatives to search in finding information, such as orienteering strategies [3], the effect we observed appears to be a difference in strategy for making sense of information, and some users appeared to have a strong preference for one strategy over the other.

Usability Results

Searching. Our tools provide several ways to initiate and navigate through a search. We hypothesized these user preferences:

1. For initiating search at the corpus level: a preference for clicking in the notebook, as this requires no typing.
2. For initiating within a document: automatic search seeded by the corpus-level query, as this requires no additional typing, once the corpus-level search is run.
3. For navigation, a preference for the arrow keys to open query-matching documents in sequence, as this takes less effort than finding an icon per document and clicking on it.

In fact, we did not find any of these hypothesized preferences. Some users did initiate a search by clicking on an entity in the notebook. However, other users initiated most queries by typing in the search boxes. In some cases,

users would click on an entity in the notebook as a quick way to get its string into the search box but would then modify the query by typing. Furthermore, users who largely initiated queries by typing were inconvenienced by the current design which replaces the query with a new one as a side effect of selecting an entity in the notebook.

Users had a similar reaction to the automatic within-document search. Some users would provide a broad query to find documents but then want to find a more restricted set of phrases within the document. Again, these users were inconvenienced when the system overwrote their manually-written query with a computed one.

To track use of documents found by a search, users happily referred to the icon array, and found the next document by position in the array. Users made little or no use of arrow keys to read through documents in order of search score.

Reading. Reactions to automatic highlighting in the documents were mixed. Some users found them helpful or at least easy to ignore when they were not wanted. Other users found them so distracting they wanted to turn them off. Because all entities (below a predetermined length) placed in the notebook become highlighted in the documents, any entities accidentally placed in the notebook caused highlights that were particularly counter-productive.

Note-taking. To import text, we hypothesized that users would make heavy use of the entity quick click method of importing. Quick click was indeed the greatly-preferred method and users frequently commented on the value of this function. Heavy reliance on quick click occurred even though software errors caused incorrect copying with some selection gestures. Subjects occasionally complained, however, that the notebook was too far from the document reader, causing excessive mouse travel and eye movement. More broadly, users valued being able to take notes easily, within the same context as the documents being read.

To comment on information, users typed their own information into the notebook. They usually preferred this mode to the use of snap-together statements, which were seen as more labor-intensive to construct. Comments were also added as group or section headers.

To organize information, our subjects created between 8 and 22 entity groups each, with a mean of 14.9. All users grouped these entity groups further, making between 1 and 5 supergroups, with a mean of 3.8. Almost all entities were placed in a group of some kind with a mean of only 2.3 entities ungrouped at session end. We were encouraged that no user reported a problem with managing this number of notes. The study also revealed that subjects care about the order of objects within a group; when a design flaw in snapping made it too difficult to place an entity at the top of

a group, subjects created workarounds, including placing entities outside the group but nearby, to act like titles.

CONCLUSIONS

Our study confirms that Entity Workspace can be used by novices for medical sensemaking and also raises new research questions. We found that users have different question answering strategies, like the Reader and Searcher strategies. We would like to understand when different strategies are effective or attractive for different combinations of tasks and tools.

Reactions to our tools show both enthusiasm for an integrated work environment and the need for more work on the technology supporting each sensemaking task. In particular, improved designs are needed for search tools that perform both inter-document and intra-document search, track reading progress through query results sets, and take advantage of information captured in a notebook; these new designs will probably need to be simpler and more unified than our current design. Users will want additional control over automatic highlighting of content, both to ensure relevance and to accommodate users with greater or lesser tolerance for highlighting in general. The heavy use of entity quick click in the study suggests that further improvements will be valuable, including ways to reduce eye movement and mouse travel and to recover from unintended copying. Likewise, feedback on grouping suggests that group titles will be valuable, along with more precise ways to order elements within groups.

This research was funded in part by the Advanced Research and Development Activity NIMD and ARIVA programs (MDA904-03-C-0404).

REFERENCES

1. Bier, E., Ishak, E., and Chi, E. Entity quick click: rapid text copying based on automatic entity extraction. *Extended Abstracts of the Conference on Human Factors in Computing Systems (CHI 2006)*.
2. Bier, E., Ishak, E., and Chi, E. Entity Workspace: an evidence file that aids memory, inference, and reading. *Proceedings of Intelligence and Security Informatics (ISI 2006)*, pp. 466-472, Springer-Verlag.
3. Teevan, J., Alvarado, C., Ackerman, M., and Karger, D. The perfect search engine is not enough: a study of orienteering behavior in directed search. *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2004)*, pp. 415-422.
4. Wright, W., Schroh, D., Proulx, P., Skaburskis, A., and Cort, B. The sandbox for analysis: concepts and evaluation. *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2006)*, pp. 801-810.