

Sensemaking for the rest of us

Daniel M. Russell, Robin Jeffries, Lilly Irani

Google

1600 Amphitheatre Parkway
Mountain View, CA 94304
{drussell, Jeffries}@google.com
lilly.irani@gmail.com

ABSTRACT

Sensemaking is simple—it’s the way people go about their process of collecting, organizing and creating representations of complex information sets, all centered around some problem they need to understand. While there are a number of tools and systems to support sensemakers, the largest group of users—people who are not highly technical—might be left out of the sensemaking tool space by our field’s focus on extremely high-end visualizations and tools. In a series of studies we looked at the effect of tools on different sensemaking processes, finding that for many sensemaking jobs, complex tools and advanced representational constructs don’t actually seem to help, but can hinder the process of sensemaking.

Author Keywords

Sensemaking, notetaking, notebook, spreadsheets

INTRODUCTION

Most people don’t actually need to visualize high dimensional data sets that incorporate time-varying series data. Instead, it’s pretty clear that a large fraction of everyday sensemaking tasks are fairly simple: one collects a bunch of data on a topic of interest, and then needs to lightly re-organize it in order to determine the best option among many to take. This description can apply to everything from selecting a new (or used) car to buy, to deciding which grad school to attend, to understanding where to go on vacation.

But a great deal of sensemaking work focuses on high-end sensemakers. That is, people with vast amounts of data and intricate, fundamentally complicated problems.

What about the rest of us? The just plain folks that need to figure out how to organize a dinner party for twelve people?

Submitted to the CHI 2008 Workshop on Sensemaking,
Florence, Italy. April 7, 2008.

NOTETAKING

Foraging and collecting

Sensemaking most often begins with collecting notes on a topic. The collecting and foraging phase might be of quite long duration, but as is apparent from earlier studies of sensemaking [Card, 2004; Pirolli, 2007], sensemaking tasks of any interesting complexity usually build upon a base set of collected notes and documents.

Collected notes, alas, may be in many formats, shapes and sizes. In traditional intelligence analysis, the “shoebox” is the collection point for notes on a topic. [Wright, 2005] Many systems have tried to emulate this practice with digital collections of clippings, documents and user-generated content. [Halasz, 1988]

As in all notetaking practice, the supposition is that the collection will be useful. Individual practices differ tremendously, but collections of materials are often created, if only for the reminding value of the items. An important question is, once collected, what does one do with the notes?

STUDY 1—NOTETAKING PRACTICE

In late 2006, we conducted a field study of the ways Google’s Notebook product was being used by people in their own tasks.

Background: Google Notebook is a browser-plugin, web-based notebook service that can be used to create “notebooks.” Each notebook is a collection of “sections,” each of which can hold text, links or images clipped from web pages (with the backlink to the source material automatically maintained), or it can also contain user-written text. Because they’re made available through a browser plug-in extension, the notebook is always available to the user when web browsing. Google search results pages also have a “Note this” link that automatically creates an entry for the displayed result in the currently selected notebook. (See Figure D at the end for an example of a Google Notebook.)

Study: We obtained permission to examine 163 randomly selected notebooks from Google Notebook users in North America. The notebooks were provided by 107 unique contributors (some shared more than one notebook with us).

Of these notebooks, we found that 135 were private (i.e., not shared), 25 were shared with from 1 to 8 collaborators, and 13 were public.

Of these 163 notebooks, we found that the majority were single-topic notebooks (see Table 1). Nearly as large in number were “other” notebooks, mostly non-topical collections of clippings from web pages and user-generated text.

39	notes on a particular topic
35	other (special, one-off note collections, not focused on a single topic)
22	bookmark collection
18	random notes (no discernible topic)
15	code (or programming)
13	to do list
12	travel notes and planning
9	shopping notes

Table 1: Notebook topics in study set

The topics were wide ranging (although admittedly, many had a clear technology focus, suggesting that the sample had an early-adopter bias). Notebook topics included: film links, Social Bookmark Sites, misc, Mapping and Navigation, Linux Tools, Python, Palm, Mac OS X, Contests, Good Songs, Pub Night Ideas, To Dos.

During January of 2007 we intensively studied these notebooks, looking for typical use patterns. We noted how often the notebooks had been used, how often sections had been created and reorganized.

Hypothesis: Our initial hypothesis was that people would use their Google Notebooks in ways that were similar to other, ordinary collections of materials—that is, that they would collect broadly, and then sort, organize, cull, and collate interesting materials together.

Analysis: For each of the notebooks we counted how often each had multiple labeled sections, how often content had been deleted from a notebook, and how often each notebook had sections moved or renamed.

To our surprise, only 48% of the notebooks had sections—the rest were just one long, often undifferentiated section in a single notebook.

Perhaps even more surprising, only 5% of these notebooks had ever seen a section moved from one place to another. Roughly the same number of notebooks had ever had a section deleted.

What was going on? Apparently, notebook users were just collecting and then never deleting or restructuring the

content of their notes. This was not what we’d expected at all. Instead, we found a great deal of list creation (essentially just lists of notes, great movies or to-do items) with little sensemaking behavior.

Field study: To get a sense of what was happening, we conducted 5 in-depth interviews with owners of the most used (and probably most interesting) notebooks.

In late January, 2007 we conducted remote interviews by telephone, using a screen-sharing system to connect us with our distant users. We asked each to give us a tour of their notebooks and to describe how they used them, with both positive and negative examples of notebook use. Each participant was compensated for their participation.

It quickly became clear that notebook users were not doing quite what we’d expected. Instead of a notebook *sensemaking* system, Google Notebook was being used as a clever way to collect notes in a list, and then manage focus of attention for mental (or otherwise non-traceable) activity.

Did Google Notetakers do Sensemaking?

Google Notebook users seemed to be doing a large amount of data collection, most probably gathered while foraging through resource materials. [Pirolli, 2007] While this is a useful and important precursor to sensemaking, it isn’t sensemaking as defined in [Russell, 1993]. The data was collected, examined and read, but showed little evidence of any significant manipulation or restructuring.

With this sample set of notebooks, we were able to measure the quantity and kinds of edits to the notebooks. Much to our surprise, few of the notebooks showed any deletions (< 5%), and even fewer showed evidence of restructuring. There was a great deal of list creation, with items ranging from pure text fragments to web-page snippet and bulleted fact collections.

STUDY 1—NOTETAKING DISCUSSION

Why so little restructuring? Why so little “sensemaking” as we had defined it?

Initially, we had thought that supporting *data collection and restructuring* would lead to sensemaking behavior. But that seems not to have been the case. In fact, the literature suggests that we should have known better. Note-taking improves sensemaking (or at least success on a sense-making task), but only if the notes involve a useful representation of the task.

[Castello & Moreno, 2005] note that just transcribing information from source materials (books, oral presentations, etc.) helps a little, but re-organizing, reflecting, amplifying and synthesizing the notes seems to help in information re-use and consolidation quite a bit.

[Trafton & Trickett, 2001] note that “the less structured the notetaking environment, the better” for performance on their representative problem solving tasks. They found that

pre-imposed structure seems to damage task performance unless the structure matches the task very well.

[Melenhorst, 2004] observes that notetakers are sometimes unsure why they take particular notes—they have a general sense that they will be useful in the future, but often don't know exactly what the future need will be. That is, they very often do not filter their notes, and perhaps only weakly structure them when they're unsure of the future task.

This is a telling point: People take notes in anticipation of future need.

Finally, [Shipman & Marshall, 1999] show that formalism (especially pre-defined structures) in notetaking / argumentation systems can damage performance

Lesson: Less imposed structure is better in the absence of a clear, strong model of the task to be done with the notes.

STUDY 2—SPREADSHEETS AND REPRESENTATIONS

After our first study, we turned our attention to how people were using spreadsheets to gather and organize their data. Although less capable of gathering and representing a wide variety of data (in that they have no images or sophisticated formatting of content), spreadsheets have proven incredibly useful at collecting and organizing a wide variety of information for a wide variety of tasks. [Nardi, 1990]

In this study we were hoping to find evidence that users would make sense of their collected data; that is, we would see spreadsheet structure and data collected and transformed to satisfy an analysis task need.

Background: Google Spreadsheet is, like Google Notebook, a web-app providing basic spreadsheet capabilities to users over the internet. The data is stored in the cloud, and can be created, accessed and shared with other users through a straight-forward web browser interface.

Spreadsheet Study

In this somewhat smaller study, we collected 24 spreadsheets from 22 people. As in the first study, the spreadsheets (with their entire edit history from inception until study date) were collected from a broad range of people covering a variety of topics. Sample topics: Group holiday planning; Team Projects people and task allocations; planning task list; onsite-summit invitation list; finding great B&B, etc.

Hypothesis: As in the Notebook case, we were searching for evidence that sensemaking was taking place in Spreadsheet user patterns. Since the Spreadsheets were not explicitly designed for collection (as the Google Notebook was), we didn't expect to see large amounts of collected information, although we did expect to see representation shifts over time as data was collected and reorganized to suit the task at hand.

Revision: 3 week(s) ago - - Pasted		« Older	Newer »
Day	3/1/2006	3/2/2006	
Date			
Day	Date		
Thursday	1-Mar-2006		
Friday	2-Mar-2006		
Saturday	3-Mar-2006		
Sunday	4-Mar-2006		
		Allen School District Spring Break	
Monday	5-Mar-2006		
Tuesday	6-Mar-2006		
Wednesday	7-Mar-2006		
Thursday	8-Mar-2006		
Friday	9-Mar-2006		
Saturday	10-Mar-2006		

Figure A: A sample spreadsheet, 3 weeks before the study period.

Analysis: We examined the 24 spreadsheets, analyzing each of their changes over their lifespan, looking for any “significant structural change.” We defined this as an edit or change to the spreadsheet that was clearly an attempt to alter the representational properties of the spreadsheet.

These changes included: adding a new row or column with a new kind of data; re-factoring data into different rows or columns to expose some new aspect of the data; adding a new representation mechanism (such as adding color formatting to create bar charts out of cells); merging row or columns together to reduce complexity of combine aspects of the data; rotating the data from row-order to column-order. In all, we found these 5 kinds of changes to cover nearly all of the representational changes in our test set.

Trivial changes (such as font changes, text color changes, adding a value into an existing structure, etc.) were not counted.

Somewhat surprisingly, 19 of the spreadsheets were used primarily in a non-numeric way (such as Figures A and B), suggesting that the spreadsheets are often being used to collect and coordinate data, rather than purely for computational work.

More intriguingly, as we examined the edit history of each of the spreadsheets, 18 of them (75%) have significant structural changes over the course of their lifetimes. The spreadsheets had lifetimes from 2 weeks to 1 year before we collected them for our study.

Spring Break Planning Autosaved at Dec 12, 2006 8:08:18 AM PST

File Edit Sort Formulas Revisions

Browse Revisions

Revision: 3 week(s) ago - dminogue - Pasted < Older Newer >

Day	Friday	Saturday	Sunday	Monday
Date	3/2/2007	3/3/2007	3/4/2007	3/5/2007
Allen S.D.				
Newton S.D.				
Amanugi				
Beardsley				
Bell				
Boland				
Minogue				
Mitori				

Day	Friday	Saturday	Sunday	Monday
Date	4/6/2007	4/7/2007	4/8/2007	4/9/2007
Allen S.D.	Good Friday			Bad Weather Day
Newton S.D.	Good Friday			

Figure B: The same spreadsheet as in Figure A, a few days later. Note the important structural changes that affect the spreadsheet’s ability to represent aspects of the problem (scheduling).

Did Google Spreadsheet Users do Sensemaking?

As became quickly clear, Google Spreadsheet users in fact DO a significant amount of representation structural changes and modification to fit to the task.

While significant structural changes seemed to occur in bursts of activity, we did not have sufficient data to warrant any kind of detailed analysis of frequency of edits during the course of the spreadsheet.

However, we did notice a kind of repeating theme of making structural edits, followed by a period of testing before making another structural edit. This fits with the pattern described in [Russell, 1993] and matches our intuition that big changes happen, then are tested to determine if the new representation can usefully solve the problem.

STUDY 2—SPREADSHEETS DISCUSSION

It’s clear that sensemaking behavior is taking place in the hands of spreadsheet users. When we reviewed the edit history of each spreadsheet with its users, it became clear that the structural changes are driven by clearly perceived task needs, and not an abstract notion of better fit to task. That is, task requirements clearly drove changes in the representation directly on a moment-by-moment basis. This has caused us to modify the original sensemaking model diagram from [Russell, 1993] to add two components that deserve their own recognition as salient phases. (As shown in Figure C.) Foraging / collection really does seem to be a separable phase, as suggested by our spreadsheet and notebook users who would overtly collect information to be organized in a later step. And we have also added a backlink between the task structure and the search for a representation. Representation design and

evolution is not just a drive to an increasingly better representational fit, but is a separate evaluation step that can be satisfied as needed.

As we saw in many cases, the evolution of the spreadsheet representation was rapid and direct, even if some of the transformations were painfully laborious (such as transposing rows for columns, as in Figures A to B). But at all points during the evolution, the role of the representation under construction (the spreadsheet) was a function of the overall goal. In every case, there was a clear deliberation on the part of the user—will this representational shift make the target task simpler, easier.. or even possible at all?

There have been a few research studies that illustrate the ways in which spreadsheet design reflects both the role of convention and social understanding. [Nardi, 1990] [Brown, 1987]

DISCUSSION

It’s now clear that sensemaking behavior takes place in the spreadsheets of our users, and very little (if at all) in Google Notebook users. The question is equally clear: Why? What makes the two systems so different in use?

Speed Matters: In an earlier study [Russell, 2005] we found that we could dramatically affect the way people used a sensemaking system (Grokker) by making small changes in the details of the interface. In particular, we were able to duplicate the speed and use patterns of a paper document collection in our online soft clustering system. The most important finding from that study was that small changes in speed would make a large behavioral difference.

This finding is reflected in our observations of Notebook users: as we saw during our detailed interviews, even practiced Notebook users would take around 10 seconds to get to the note they wanted to lookup or edit. Since this is true for almost all actions, this turns out to be uncomfortably long, especially since the Grokker study showed that changes of less than one second in action time would affect user behavior. In a computational environment where experts frequently switch tasks in less than 3 seconds (e.g., alternating from web browsing to using a text editor), a subtask that takes 10 additional seconds is a forbidding interruption.

Further, we found that users could not restructure notes in a Notebook by dragging items in the smaller “toast” view. (Called “toast” because it pops up, like toast, in an overlay on the browser window.) Restructuring requires being in the “full page” view, an operation that would take an average of 4 seconds. But what was probably worse was that once the user opened the Notebook into the full page view, there was no “minimize” operation to get back to the original state. The only thing a user could do would be to close the window (or tab) and re-open the Notebook from scratch. Although state was preserved, the entire operation felt clunky, and required considerable time.

Search Matters: Another confounding behavior of Notebook is that the search mechanism behaved in an odd, often unpredictable fashion. Find-on-page (using Control-F) would frequently make the found selection disappear from view, leaving the user to wonder what had happened. Was the string found or not?

An even more egregious confusion was the problem of cross-notebook search. While Notebook has a search box at the top of each full-page view, the search terms were searched as tokens, not as strings. That is, if the user looks for a note with the string “point,” the search would NOT find notes with the terms Northpoint or Southpoint, as the search operates only on full terms.

Clearly, if one is taking notes for the purpose of future reference, a weak find mechanism is not a help. This is particularly problematic when trying to find a common thread or structure in a set of notes.

Navigation Matters: Lastly, when using the Notebook, after doing a search or a navigation move to another notebook, there was no way to go back to the previous notebook by a keyboard shortcut or click on the Back icon. Similarly, if one remembered to toggle between notebooks A and B, when the user clicked back to notebook A, the viewing position would be lost. This made working with notebooks longer than one screen-full very problematic. Every cross-notebook analysis would take multiple clicks just to do comparisons.

All in all, the implementation details of Notebook worked against allowing restructuring and notetaking behavior that required anything more than fairly straightforward note collection.

Simple structuring: In Google Spreadsheets, the canonical spreadsheet model is clearly supported and widely understood. If anything, the Google version of spreadsheets is an even simpler version of the widely used (and well-learned) Microsoft Excel model.

Thus, when Spreadsheet users would move, copy, and restructure their spreadsheets, there was very little to stand in their way. Well practiced behaviors would do virtually the same thing in Google Spreadsheets as they had learned elsewhere. The simple structuring tools of spreadsheets are easy to build, simple to communicate to co-users and integrate smoothly with other tools.

One has the sense that users are constantly wrangling with the Google Notebook structuring tools, whereas with Google Spreadsheets, they tend to not think much about the tool, but focus their attention on the problem, its representation and its solution.

It’s also true that people have a great deal of practice, both as individuals and as members of groups that share spreadsheets. There is a commonly accepted practice, and continual reinforcement through use. By contrast, Notebook is sufficiently different than other notetaking

applications that common practice and social reinforcement effects do not have time to take.

CONCLUSION

The implications of these studies are clear: sensemaking tools must fit into the flow and style of a user’s work practices. They should build on patterns of representations that are well-understood, and both fast and simple to use. Tools that are used rarely, or that have a variant working style need to accommodate the sensemaking behaviors that are common to the people that will be using the tool.

The current style of sensemaking tools is often to design high-end, complex, sophisticated tools. The intent is to support high-end, sophisticated users and to derive lessons from the behavior of extreme users.

However, the “rest of us” is an enormous audience, well worth attention. As we have seen, the design problem for sensemaking tools—even ones as simple as note taking tools—are exquisitely sensitive to the design of their interfaces. Careful attention to systems that provide simple, integrated and facile tools is the critical path.

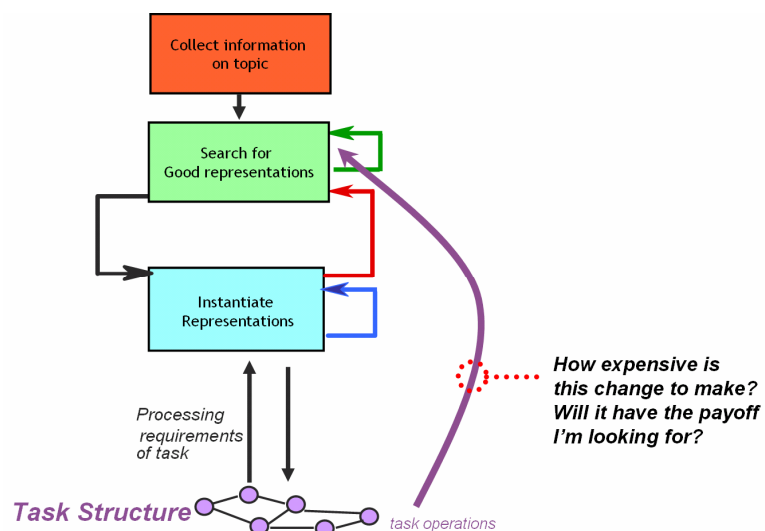


Figure C: The sensemaking loop revised to include an explicit collection/foraging step at the outset, and another feedback loop that evaluates the tradeoff between making a representation change versus accepting a less good task performance.

ACKNOWLEDGMENTS

Thanks to our participants (who are many and anonymous). Our gratitude also goes to many of the people at Google who make the products and help us do the analysis. Unfortunately, they are too many to number here. But we are thankful for you all.

REFERENCES

1. Beale, R. Supporting serendipity: Using ambient intelligence to augment user exploration for data mining and web browsing, *International Journal of Human-Computer Studies*, 65(5), p 421-433, (May 2007)
2. Brown, P. S., Gould, J.D. An experimental study of people creating spreadsheets, *Transactions on Office Information Systems* 5, 258-272 (1987)
3. Card, S. K., Hong, L., Mackinlay, J. D., Chi, E. H. 3Book: a scalable 3D virtual book, CHI 2004, p 1095-1098, Vienna, Austria. (May, 2004)
4. Castelló, M. and Monereo, C. Students' Note-Taking as a Knowledge-Construction Tool. L1-Educational Studies in Language and Literature. Vol 5, pp. 265-285. (2005)
5. Halasz, F. G., Reflections on NoteCards: seven issues for the next generation of hypermedia systems, *Communications of the ACM*, v.31 n.7, p.836-852, July 1988
6. Melenhorst, M. Supporting professional readers of online documents. Proc. of ACM CHI 2004 Conference on Human Factors in Computing Systems, v.2 p.1059-1060, (2004)
7. Nardi, B., Miller, J. An ethnographic study of distributed problem solving in spreadsheet development. CSCW 1990, 197-208. (1990)
8. Nardi, B., Miller, J. Twinkling lights and nested loops: Distributed problem solving and spreadsheet development. *International Journal of Man-Machine Studies*, 34, 161-184 (1991)
9. Pirolli, P. *Information Foraging Theory: Adaptive Interaction with Information*, Oxford (2007)
10. Russell, D. M., Stefik, M. J., Pirolli, P., Card, S. K. 1993. The Cost Structure of Sensemaking. Proc. InterCHI '93, Amsterdam, Netherlands, 269-276, (1993)
11. Russell, D. M., Slaney, M., Qu, Y., Houston, M. Being literate with large document collections: Observational studies and cost structure tradeoffs, HICSS 2006, Hawai'i (Jan, 2006)
12. Reimer, Y. J., Brimhall, E., and Sherve, L. A study of student notetaking and software design implications. Proceedings of the 5th IASTED international conference on Web-based education. Puerto Vallarta, Mexico, pp. 189 – 195, (2006)
13. Shipman, F., Marshall, C., Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems, Proc. of Computer-Supported Cooperative Work (CSCW), (1999)
14. Trafton, J. G. and Trickett, S. B. Note-Taking as a Strategy for Self-Explanation and Problem Solving. *Human-Computer Interaction*, v.16. p.1-38 (2001)
15. Wright, W., Schroh, D., Proulx, P, Skakrskis, A., Cort, B., Advances in nSpace – The Sandbox for Analysis, Conference on Intelligence Analysis (2005)



SketchUp Home

Developed for the conceptual stages of design, this powerful yet easy-to-learn software allows for quick and easy 3D form creation, the result is an ...
www.sketchup.com/ - 16k - [Cached](#) - [Similar pages](#) - [Note this](#)

- [Pro Download](#)
- [Google SketchUp](#)
- [Google SketchUp Pro 6 Product](#)
- [Support Plugins](#)
- [Google SketchUp Pro Overview](#)

[More results from sketchup.com »](#)

SketchUp Google Plugin

There are now two **SketchUp** exporters for KMZ and a KMZ4 for use with Google Earth 4. ...
www.sketchup.com/index.php?id=408 - 19k - [Cached](#)

Google SketchUp - Home

A 3D software tool that combines a tool-set with a place models using real-world coordinates and share ...
sketchup.google.com/ - 9k - [Cached](#) - [Similar pages](#)

[Google SketchUp - Download](#)

Sponsored Links

Google SketchUp (Free)

Create, view, & modify 3D models. Download your own copy, free.
sketchup.google.com

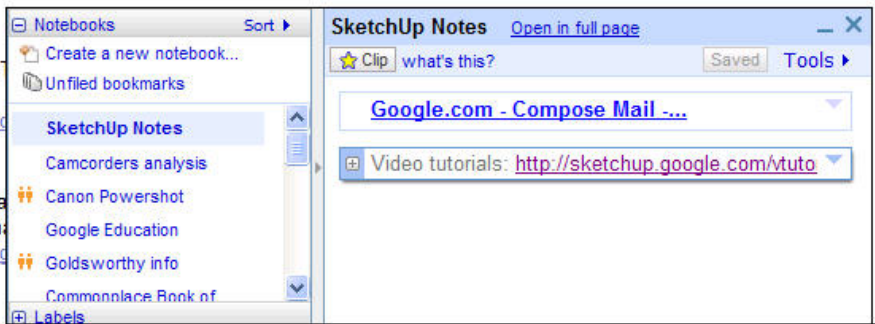


Figure D: A Google Notebook can be used as a popup (seen here) that appears in the context of the websearch (or other web-based information gathering task). The notebook appears as an inset in the lower right of the browser. Each notebook consists of sections, each of which can be expanded to view a complete note, or closed (as shown here) to show only the title of the section. Sections can be dragged around to reorder as needed. Different notebooks are shown in the left-hand-side pane. Shared notebooks are denoted with the golden-yellow two-person icon.