

The ScratchPad: Sensemaking Support for the Web

David Gotz
 IBM T. J. Watson Research Center
 19 Skyline Drive, Hawthorne, New York 10532 USA
 dgotz@us.ibm.com

ABSTRACT

The World Wide Web is a powerful platform for a wide range of information tasks. Dramatic advances in technology, such as improved search capabilities and the AJAX application model, have enabled entirely new web-based applications and usage patterns, making many tasks easier to perform than ever before. However, few tools have been developed to assist with sensemaking tasks: complex research behaviors in which users gather and comprehend information from many sources to answer potentially vague, non-procedural questions. Sensemaking tasks are common and include, for example, researching vacation destinations or deciding how to invest. This paper presents the *ScratchPad*, an extension to the standard browser interface that is designed to capture, organize, and exploit the information discovered while performing a sensemaking task.

Categories and Subject Descriptors: H.5.2 [Information Systems]: User Interfaces

General Terms: Algorithms, Design, Human Factors

Keywords: Sensemaking, Visual Analytics, Web Browser, WWW

1. INTRODUCTION

Today's web tools are designed to efficiently support *browsing* and *information seeking* tasks. Web users can use their browsers to quickly locate train schedules or retail locations, leaf through digital product catalogs, or browse blogs and videos via large collections assembled via crawling, tagging, or other methods. Additionally, users can perform *procedural* tasks, such as purchasing a book, paying bills, or getting directions, directly from their web browser.

However, there is a class of behaviors, *sensemaking tasks*, which is not well supported with today's tools. Sensemaking is a complex process where users gather, comprehend, and connect nuggets of information from many sources to answer potentially vague, non-procedural questions. For example, business intelligence and exploring vacation destinations are both sensemaking tasks.

Such information tasks demand more than the efficient information retrieval and presentation offered by today's browsers. They require additional tools that enable the marshaling of discovered data, the expression of new insights,

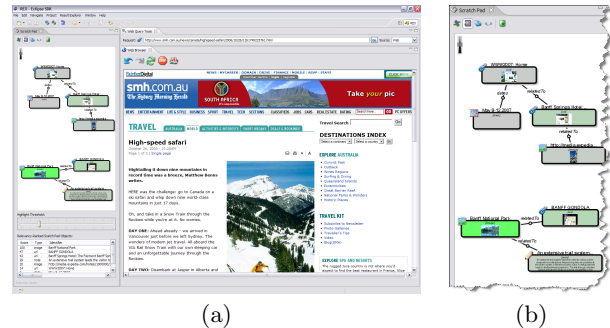


Figure 1: (a) A ScratchPad-enhanced web browser and (b) a close-up of the ScratchPad sidebar.

and the exploitation of previously synthesized knowledge in future stages of analysis.

2. THE SCRATCHPAD

The *ScratchPad*, shown in Figure 1, is an extension to the standard browser interface designed specifically to assist in sensemaking tasks. The ScratchPad provides tools for users to capture both user-created insight as well as information found on the web while completing a sensemaking task. Captured information is visually displayed to the user for their own review and re-access. More importantly, the captured information is exploited algorithmically to help users “connect the dots” as they go about their research activities.

The ScratchPad offers two primary contributions. First, it provides a set of visual interaction tools which allow a user to easily capture, represent, and manipulate information at various granularities. Second, it defines both an algorithm and user interface technique for calculating and conveying the relevance of previously captured information to a user's current browsing behavior.

2.1 Visual Presentation and Interaction

The ScratchPad is designed to store, organize, and display the information collected over the course of a sensemaking task. The ScratchPad is presented to the user as a re-sizable vertical sidebar to the left of the main browser display area. The side-by-side display of the ScratchPad and browser allows users to quickly move between the two spaces. Users of the ScratchPad have access to several tools for interactively creating and manipulating objects within the workspace.

Capturing Information. The first method for creating a new object in the ScratchPad is *snapshotting*. Snapshots record a page's URL together with (1) a graphical image of




Figure 2: (a) Snapshots are created by dragging the link icon from the URL bar to the ScratchPad. (b) Data at finer granularities can be captured by dragging specific page elements, such as snippets of text.

the browser’s rendering of the page, and (2) a copy of the page’s textual content. Snapshots are created by dragging an icon from the URL bar above the browser and dropping it on the ScratchPad (see Figure 2[a]). Once captured, the snapshot is displayed graphically, similar to other bookmarking tools [1, 3]. Alternatively, users can collect information at a much finer granularity using “drag-and-drop” from a web page. This technique can capture individual links, images, or text fragments (see Figure 2[b]).

Users are also able to create new objects in the ScratchPad from scratch, without selecting and dragging anything from the browser. A toolbar at the top of the ScratchPad allows users to create arbitrary new objects using a dialog box. These objects can be used to help organize information captured from the browser (e.g. creating folders), or to express insight created by the users themselves as the sensemaking task evolves (e.g. text notes).

Manipulating Information. Users are provided with two tools for the manipulation of ScratchPad objects. First, objects can be *modified* by altering stored properties. The modification capability of the ScratchPad is roughly analogous to the functionality of traditional bookmark organization tools which allow you to modify titles and contents of bookmarked items.

Second, objects can be *linked* (and unlinked) to express relationships to other objects on the ScratchPad. For example, a user might decide that two web pages are related to each other, or that a particular text note is related to a captured image. Links are created by a simple drag-and-drop gesture. Implicit links between objects are also created based on the set of folders used to organize the ScratchPad.

Links allow users to record relationships between entities in the ScratchPad as they are discovered, the same way they might using pencil and paper. This behavior is patterned after the hand-written notes shown in Figure 3 which show actual notes observed during a user study exploring sensemaking behavior [2]. Links are also exploited by a relevance detection algorithm that helps detect overlooked connections between a user’s ScratchPad notes and the current web page being displayed in the browser.

2.2 Information Relevance

One of the key features of the ScratchPad is *relevance detection*. As a user browses the web in search of new information related to their task, the ScratchPad employs an algorithm that compares the information stored in the tool’s graph-based data model with the content of the web browser. As connections are found, a numerical measure of the relevance is computed and conveyed visually to the user.

The ScratchPad provides two mechanisms for displaying the estimated relevance values. First, parameterized object highlighting within the graphical ScratchPad display enables

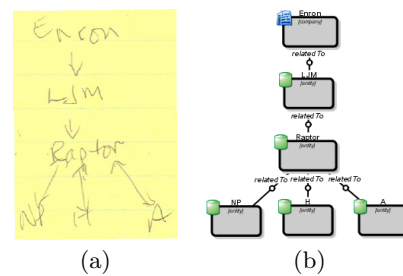


Figure 3: (a) Hand-drawn notes from a user study examining analyst behavior during sensemaking tasks. (b) The ScratchPad’s interface design was motivated in part by the study’s findings.

qualitative, “at-a-glance” detection of relevant regions of the context graph. Second, a sortable table presentation of the of the ScratchPad objects allows users to examine the relevance score assigned to each node quantitatively.

Parameterized Object Highlighting. The ScratchPad interface displays the relevance score assigned to each node by highlighting the graphical objects in various shades of green. Stronger green tones indicate qualitatively that a node has a stronger relevance to the content displayed within the web browser. In contrast, less intense highlights indicate a weaker relevance score. The highlight intensity is updated dynamically each time a user browses from one web page to the next.

The highlighting behavior is evident in Figure 1, where a user is reading a web page about winter sports in the Canadian rockies as part of their research task exploring a trip to Banff for the WWW 2007 conference. The ScratchPad objects scored as most relevant deal with Banff National Park, a gondola attraction, and some hiking trails. Also shaded green is a node for the Fairmont Banff Springs Hotel. Just as importantly, nodes corresponding to the conference web page and dates are not highlighted, indicating at a glance which portion of the user’s previous work is most relevant to the page currently being viewed.

Tabular Relevance Presentation. The ScratchPad also provides a quantitative presentation of the relevance data to complement the qualitative nature of the highlight-based graphical display. Below the graphical note-taking area, the ScratchPad provides a tabular presentation of the relevance data.

3. FUTURE WORK

The ScratchPad is only an initial prototype and several areas would benefit from additional research. In particular, improving the selectivity of the relevance metric, perhaps by incorporating text analytics, would make the visual relevance presentation even more effective. In addition, user studies to examine the value of such a tool are required.

4. REFERENCES

- [1] B. Amento, L. G. Terveen, W. C. Hill, and D. Hix. Topicshop: enhanced support for evaluating and organizing collections of web sites. In *UIST*, pages 201–209, 2000.
- [2] D. Gotz, M. X. Zhou, and Z. Wen. A study of information gathering and result processing in intelligence analysis. In *IUI 2006 Workshop on IUI for Intelligence Analysis*, 2006.
- [3] N. Jhaveri and K.-J. Raiha. The advantages of a cross-session web workspace. In *CHI Extended Abstracts*, pages 1949–1952, New York, NY, USA, 2005. ACM Press.