Rich-Prospect Browsing Interfaces

Stan Ruecker
University of Alberta, Canada

INTRODUCTION

Everyone who has browsed the Internet is familiar with the problems involved in finding what they want. From the novice to the most sophisticated user, the challenge is the same: how to identify quickly and reliably the precise Web sites or other documents they seek from within an ever-growing collection of several billion possibilities?

This is not a new problem. *Vannevar Bush,* the successful Director of the Office of Scientific Research and Development, which included the Manhattan project, made a famous public call in *The Atlantic Monthly* in 1945 for the scientific community in peacetime to continue pursuing the style of fruitful collaboration they had experienced during the war (Bush, 1945). Bush advocated this approach to address the central difficulty posed by the proliferation of information beyond what could be managed by any single expert using contemporary methods of document management and retrieval.

Bush's vision is often cited as one of the early visions of the World Wide Web, with professional navigators trailblazing paths through the literature and leaving sets of linked documents behind them for others to follow. Sixty years later, we have the professional indexers behind *Google,* providing the rest of us with a magic window into the data. We can type a keyword or two, pause for reflection, then hit the "I'm feeling lucky" button and see what happens.

Technically, even though it often runs in a browser, this task is "*information retrieval*." One of its fundamental tenets is that the user cannot manage the data and needs to be guided and protected through the maze by a variety of information hierarchies, taxonomies, indexes, and keywords. Information retrieval is a complex research domain. The *Association for Computing Machinery*, arguably the largest professional organization for academic computing scientists, sponsors a periodic contest in information retrieval, where teams compete to see who has the most effective algorithms. The contest organizers choose or create a document collection, such as a set of a hundred thousand newspaper articles in English, and contestants demonstrate their software's ability to find the most documents most accurately. Two of the measures are precision and recall: both of these are ratios, and they pull in opposite directions. *Precision* is the ratio of the number of documents that have been correctly identified out of the number of documents returned by the search. *Recall* is the ratio of the number of documents that have been retrieved out of the total number in the collection that should have been retrieved. It is therefore possible to get 100% on precision—just retrieve one document precisely on topic. However, the corresponding recall score would be a disaster. Similarly, an algorithm can score 100% on recall just by retrieving all the documents in the collection. Again, the related precision score would be abysmal.

Fortunately, information retrieval is not the only technology available. For collections that only contain thousands of entries, there is no reason why people should not be allowed to simply browse the entire contents, rather than being limited to carrying out searches. Certainly, *retrieval* can be part of browsing—the two technologies are not mutually exclusive. However, by embedding retrieval within browsing the user gains a significant number of perceptual advantages and new opportunities for actions.

The history of information *visualization* is long and interesting (e.g., Tufte, 1990). A wide range of browsing interfaces have also been developed, both as experimental prototypes and in working versions. For example, Small (1996) created a three-dimensional space containing the text of Shakespeare's plays, which readers could examine by flying among the columns. Shneiderman, Kang, Kules, Plaisant, Rose, and Rucheir (2002) designed a system for organizing photo

collections by displaying thumbnails of every photo, combined with tools for manipulating the display. Variations of Shneiderman's approach have proliferated in the last few years with the release of the application programming interface (API) for Flickr, which allows researchers to create experimental interfaces to a massive online collection of digital photos. Further examples include Bumgardner (2005), who provides a search engine based on a colour wheel, Klingemann (2006), who allows users to browse related photo tags set on the circumference of a circle with access to the images through the central text, and Turner (2006), who presents a similar arrangement of tags but includes a set of images in the centre. Some of these can be considered "rich-prospect browsing interfaces," in that they meet the following five criteria:

a. Some meaningful representation of every item in the collection is an intrinsic part of the default interface.
b. There are tools available to allow the user to manipulate the display, for instance, by reorganizing, grouping, or subsetting the meaningful representations.
c. Where possible, these tools draw on the information available in the collection, that is, a complex set of data has more and better tools than a simple set.
d. Where possible, there is more than one meaningful representation of every item in the collection, and changing the display between representations is under the control of the user.
e. The representations become the means of accessing further data. Not only can the user see what is available, but it is also made immediately accessible (Ruecker & Liepert, 2004).

Each of these five criteria will be expanded and clarified in the discussion that follows.

The philosophical assumption behind rich-prospect browsing interfaces is that people have a tremendous capacity to manage information. We are ecologically positioned to be good at looking at complex displays, such as may be found in the natural environment, and extracting what is important to us, such as whether or not that apple is ripe enough yet to eat. We have very large and sophisticated eyes, and in our brains there are multiple parallel visual systems for discriminating environmental features such as luminosity, motion, and colour. Rich-prospect browsing interfaces are an attempt to leverage these ecological advantages for the benefit of people surfing the Web (Figure 1).

Figure 1. This rich-prospect browsing interface allows people to identify pills by visually grouping pill photos. It was developed at the University of Alberta by Stan Ruecker, Lisa Given, Bess Sadler, Andrea Ruskin, and Mike Plouffe. Pill information and images were generously provided by www.drugs.com.