What Eye Movements Have Told us About User Navigation of the Web

Duncan P. Brumby (brumby@cs.drexel.edu)
Computer Science, Drexel University
3141 Chestnut St., Philadelphia, PA 19104 USA

Andrew Howes (HowesA@manchester.ac.uk)
School of Informatics, University of Manchester
PO Box 88, Sackville St., Manchester M60 1QD UK

Introduction

How do people search a web page for a link that is relevant to the achievement of their goal? Recently this question has received significant theoretical attention.

Information foraging theory (Pirolli & Card, 1999; Pirolli, 2005) has had a seminal contribution in building our understanding of how people search the web. The theory assumes that during web-based information gathering activities people are sensitive to the rate of information gain in relation to the cost of interaction. Browsing actions are assumed to be determined by the relation of navigation cues (information scent) to the user's information goal. However, estimates of information scent must also be embedded within a strategy for controlling search. A number of models of the cognitive processes that might be involved in controlling search have been proposed (Brumby & Howes, 2004; Cox & Young, 2004; Miller & Remington, 2004; Pirolli & Fu, 2003; Young, 1998).

We (Brumby, 2005; Brumby & Howes, submitted) have conducted a number of studies to evaluate some of the theoretical claims regarding how people search a web page. In these studies, tracking peoples eye movements allowed conclusions to be made about the strategies people deploy to choose between assessment and selection while searching a novel web page.

Before providing a brief overview of the main empirical findings, we will argue for the importance of adopting an eye-tracking methodology to understand how people search the web.

Why look at eye movements to understand search?

There have been a number of studies that have investigated how people locate goal relevant information on the web, most of which have relied on analysis of user logs. For example, analysis of link selection data from search engine user logs (Jansen & Pooch, 2000; Silverstein, Henzinger, Marais, & Moricz, 1998) has found that people tend to select very few links over the duration of a search session and rarely go beyond the first page of the results list. However, this methodology leaves many open questions. For instance, it is not clear whether people exhaustively assess each of the items on a page, or simply select the first goal relevant item that is encountered.

Eye-tracking studies have brought significant advantages in understanding peoples' search strategies (Brumby, 2005; Granka, Joachims, & Gay, 2004). These studies provide data that go well beyond the granularity that can be gleaned from log analysis alone. For example, Granka, Joachims, and Gay (2004) conducted a study that investigated how people interact with the results page of a popular search engine. Eye-tracking was useful because it provided additional data that could not have been gained from the analysis of usage logs alone. For instance, while it was known (and expected) that users would tend to select items located within the top few positions of the results page, it was not clear whether or not they even bothered to evaluate items further down the page prior to selection. Granka, Joachims, and Gay found that the top few items on the page receive most attention, and that the average fixation time spent on an abstract drops off sharply after only the second item. This finding indicates that people are exceptionally biased, not only to selecting, but also evaluating items towards the top of the results page (the so-called *golden triangle*).

The focus of the work outlined here is on the strategies people deploy to navigate links on a web page (as opposed to the results page of a search engine). There are important differences between these two tasks. For instance, search engines are designed to return a fairly homogenous set of results that cluster around a query term with content rich abstracts, whereas labels on a web page aim to provide discriminatory navigation cues with minimal content information for users with often different goals. The most striking difference though, is that for search engines, items are explicitly ranked by their relevance to the query term (or users search goal), whereas it is often not possible to rank item by relevance on a web page. Consequently, when searching a novel web page for a link that is relevant to a search goal, people must rely on estimates of information scent and also a strategy for controlling search.

We shall provide a brief overview of a series of experiments that tracked people's eye movements in order to understand the strategies that people adopt to choose between assessment and selection while searching a simplified web page.

Experiments

Brumby (2005) presents a number of experiments that examine the strategies people use to determine when selection of an item should occur during search of a simplified web page. The experiments used an eye-tracking methodology to monitor eye movements during search. In all, experiments were conducted that manipulated the relevance (or information scent) of the labeled links in the immediate and distal choice set, and also the number of options in the immediately available choice set. A brief overview of the main empirical findings is given below.

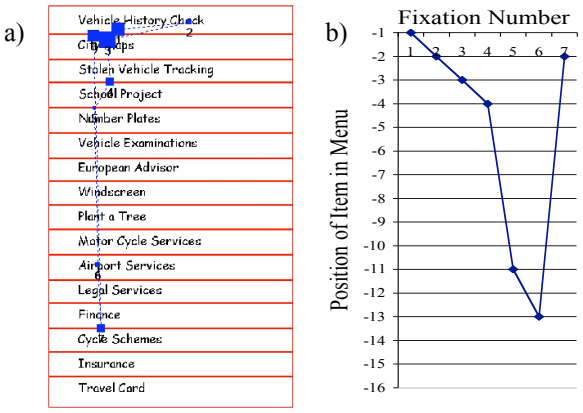

Figure 1. a) Typical eye movement trace and b) Schematic representation of the fixation sequence

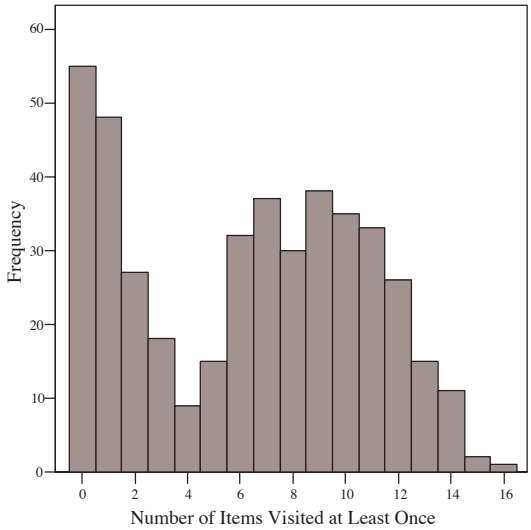

Figure 2. Distribution of the number of items visited after the initial visit to the selected item.

Overview of Method

For all of the experiments participants were required to search a simplified web page (or menu) for information relevant to a given goal statement. Participants were instructed to select labels that they believed would be likely to lead to the goal. They were also informed that within each menu there was only one correct label and that the rest of the labels were distractors. In order to impose a meaningful cost structure to this simplified laboratory task, participants did not progress to the next trial until they correctly selected the single target item from each menu. Eye tracking was performed using an ASL Pan/Tilt optics eye tracking system. For each trial we were only interested in participant's search behavior from the beginning of a trial up to the initial selection of an item.

An example menu is presented in Figure 1. The goal statement for this menu was to "Find a road map of Cardiff" and the second item in the menu "City Maps" was the target.

Overview of Main Findings

Figure 1 presents a typical eye movement trace from one of the experiments. Notice that the participant continues to check some of the remaining items in the menu after they have first looked at the target item (which they later returned to and selected).

Further analysis of the aggregate eye movement protocols gave evidence for the presence of at least two types of search behavior. Figure 2 shows the frequency distribution of the number of items visited after the initial visit to the item that was eventually selected. The distribution was clearly bimodal: There are two distributions one with a peak at 0 and the other at 9 items visited after the initial look at the selected item. This suggests that participants were sometimes choosing to select an item after visiting for the first time, a behavior we shall refer to as *first-visit-selection*. In addition, participants were sometimes choosing to visit most, but not all, of the remaining items in the menu. This behavior is reflected by the difference in the second frequency distribution in Figure 2, which approximately reflected the difference in the position of the target item.

Over a series of experiments, we found that participants search behavior was sensitive to the relevance of labels in the immediate and distal choice set and also the number of options in the immediately available choice set. In particular, participants were more likely to select an item immediately after visiting it for the first time (i.e., make a first-visit-selection) when: 1. the distractors were less relevant to the goal; 2. more of the items in the choice set had already been assessed; 3. previous experience indicated that selection was more likely to lead to success (i.e., because menu choice sets did not contain competing distractors); 4. there were fewer items in the available choice set.

It was also found that when participants located a goal relevant item and choose to continue checking the remaining items in the menu, they were more likely to skip some of these items. An experiment was conducted that manipulated the position of the target in the set and found that proportionally fewer gaze transitions were between non-neighboring items when the target item was located towards the bottom of the menu than when it was towards the top of the menu.

Conclusion

The results indicate that during interactive search people sometimes choose an option that appears good enough, but sometimes they choose to continue checking items in the menu. The results were consistent with the view that participants dynamically adjusted their evaluation of the value of further assessment (see, Brumby & Howes, 2004; Cox & Young, 2004; Young, 1998). That is, they assessed items when the time cost of assessment was judged a worthwhile sacrifice given the potential reduction in uncertainty. Participants did not make an a-priori commitment either to assessing all of the items or to assessing items until the value of the most recently assessed

item was above a threshold (what we call a simple threshold account).

Furthermore, the results suggest that people dynamically alter their search behavior depending on the relevance of the items that have been assessed. It was found that participants tended to shift from a deliberate sequential search behavior to a skipping strategy after they had located a candidate item for selection. An interpretation of the skipping behavior is that it reflects the use of a low quality, low cost assessment method during interactive search. From this perspective, it is assumed that people make choices between different assessment methods that vary in their costs and potential benefits during search.

Tracking participants eye movements was vital for these conclusions to have been made. Eye-tracking provided moment-to-moment behavioral index of users' human-computer interactions. If we assume that gaze shifts are tightly coupled with the allocation of visual perception and cognition (Findlay & Gilchrist, 2003), then the analysis of eye movement protocols allows inferences to be made regarding the sequence of items in the menu that participants choose to assess prior to selection.

References

Brumby, D.P. (2005). *An empirical investigation into strategies for guiding interactive search*. Doctoral dissertation, University of Wales: Cardiff University.

Brumby, D. P. & Howes, A. (2004). Good enough but I'll just check: Web-page search as attentional refocusing. In M. Lovett, C. Schunn, C. Lebiere, & P. Munro (Eds.), Proceedings of the *6th International Conference on Cognitive Modeling* (pp. 46-50), Mahwah, NJ: Lawrence Erlbaum.

Brumby, D.P. & Howes, A. (submitted). Strategies for Guiding Interactive Search: An Empirical Investigation into the Consequences of Label Relevance for Assessment and Selection. Submitted to *Human-Computer Interaction*.

Cox, A.L., & Young, R.M. (2004). A rational model of the effect of information scent on the exploration of menus. Poster session at the meeting of the *6th Internal Conference on Cognitive Modeling*, Pittsburgh, PA.

Findlay, J. M., & Gilchrist, I. D. (2003). Active vision: the psychology of looking and seeing. Oxford, UK: Oxford University Press.

Granka, L. A., Joachims, T., & Gay, G. (2004). Eye-tracking analysis of user behavior in WWW search. Poster session presented at the Proceedings of the *27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Sheffield, UK.

Jansen, B. J., & Pooch, U. (2000). A Review of Web Searching Studies and a Framework for Future Research. *Journal of the American Society of Information Science and Technology, 52*, 235–246.

Pirolli, P. (2005). Rational of information foraging on the Web. *Cognitive Science, 29*, 343–373.

Pirolli, P., & Fu, W-T.F. (2003). SNIF-ACT: A model of information foraging on the world wide web. In P. Brusilovsky, A. Corbett, & F. de Rosis (Eds.), *User Modeling 2003, 9th International Conference on User Modeling* (Vol. 2702, pp. 45-54), London, UK: Springer-Verlag.

Pirolli, P., & Card, S.K. (1999). Information foraging. *Psychological Review, 106*, 643–675.

Silverstein, C., Henzinger, M., Marais, H., & Moricz, M. (1999). Analysis of a Very Large Web Search Engine Query Log. *ACM SIGIR Forum, 33*, 6–12.

Young, R.M. (1998). Rational analysis of exploratory choice. In M. Oaksford, & N. Chater (Eds.). *Rational Models of Cognition*. Oxford, UK: Oxford University Press.