

Exploring How Mouse Movements Relate to Eye Movements on Web Search Results Pages

Kerry Rodden
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043, USA
kerryr@google.com

Xin Fu
School of Information and Library Science
University of North Carolina
Chapel Hill, NC 27599, USA
xfu@unc.edu

ABSTRACT

A mouse click is a proven indicator of a user's interest in a web search result. In this paper we explore the potential of a more subtle signal: mouse movements. We conducted a study where participants completed a range of tasks using Google, and we tracked both their eye movements and mouse movements. We discuss the relationship between these movements, and three different types of eye-mouse coordination patterns. We believe that mouse movements have most potential as a way to detect which results page elements the user has considered before deciding where to click.

1. INTRODUCTION

Web search engines are the source of most web users' interactions with information retrieval systems. Researchers have explored the potential of analyzing click patterns from search engines, both as a means of evaluating their ranking functions in the context of users' real information needs, and of gathering feedback to improve ranking for subsequent users (e.g. [1][8]). While extremely valuable, clicks do not tell the whole story of the user's interaction with the search results page. For example, they do not indicate *why* the user clicked on a particular result, or which other results they considered before making a choice.

A user's selection of a particular search result is based on the surrogate shown on the results page, which typically contains the page title, its URL, and a "snippet" showing one or more lines from it. The likelihood of a user clicking on a result is mostly dependent upon how promising they think it is, based on the surrogate. It would therefore be useful to have a better idea of which aspects of the surrogate users are paying attention to when making each decision about where to click. Also, in some cases it may be possible for the user to find the answer to a fact-finding question simply by reading the snippet, and many search engines now choose to present relevant information on the page directly, e.g. the definition of a word, or a stock quote. In both of these situations, no click would occur even though the user may have satisfied their information need.

Eye tracking can provide insights into users' behaviour while using the search results page, but eye tracking equipment is expensive and can only be used for studies where the user is physically present. The equipment also requires calibration, adding overhead to studies.

In contrast, the coordinates of mouse movements on a web page can be collected accurately and easily, in a way that is transparent to the user. This means that it can be used in studies involving a number of participants working simultaneously, or remotely by

client-side implementations – greatly increasing the volume and variety of data available.

Our goal in conducting this research was to investigate the potential usefulness of tracking mouse movements on a web search results page – for example, how closely do mouse movements reflect eye movements? Do people use the mouse pointer as a marker to help them read the search results, or to help them make a decision about where to click?

Unlike this study, which focused on search results pages, previous studies on the relationship between eye movements and mouse movements have been concerned either with general web pages (e.g. [2][4][10]) or with tasks that involve locating and selecting a given target item from graphical user interface menus of various lengths (e.g. [6]). It is unclear whether findings from these studies carry over to web search results. There have been several studies involving eye tracking on web search results pages, e.g. [8][9], but these have not considered mouse actions other than clicks. Researchers of implicit relevance feedback (e.g. [5][7]) have found that mouse actions on general web pages are a potentially useful signal, but they have not studied web search results pages.

2. STUDY SETUP

2.1 Tracking Mouse and Eye Movements

To capture mouse movements, we used a method similar to that described in [2] and [3]: a web proxy server inserted a reference to a piece of Javascript code at the top of every Google search results page visited. This Javascript code captured the user's mouse coordinates (and the ID of the HTML element the mouse was over) at 100 millisecond intervals, and submitted the gathered data, with timestamps, into a MySQL database every 2 seconds (or when the user left the Google search results page).

To capture eye movements, we used a Tobii 1750 eye tracker running Clearview software, with a 17-inch screen set to a resolution of 1024x768. We used IE 6 at full size on the screen. Clearview logged each URL and saved a screenshot of every web page visited during the study.

2.2 Participants

Our 32 participants (14 male and 18 female, aged 24-61) had a range of occupations and web search experience, but all were familiar with Google. 22 were from our company's user study participant database; 10 were non-technical company employees.

2.3 Tasks

We used 16 web search tasks – a sample of which are listed in Table 1. All were of the informational type, rather than

navigational or resource-related [11]. We mostly chose closed fact-finding questions with a specific correct answer, so that the participants would have a clear idea of what was required and when they were done with the task. However, in 3 of the tasks, the user had to make a decision based on their own preference.

We provided initial queries for each task (also shown in Table 1), to ensure that each user would see the same results at first. This meant that, for fact-finding tasks, we could choose queries where the answer to the question was visible on the results page itself (e.g. in the snippet of one of the results) – because we were interested to see if users would move the mouse over the answer while looking at it. For all but one of the queries, there was at least one useful result on the initial page. We did not manipulate the visible content in any way, so as well as the 10 search results, many of the pages included other elements such as sponsored links, and extra information inserted by Google.

2.4 Procedure

This was an exploratory study, not a controlled experiment, so the participants each did the same 16 search tasks, in the same order. They started from a study home page that contained one link per task. Each of these links led to a fake Google home page, with the initial query pre-filled in the search box, and the task description inserted underneath. Participants were instructed to press the “Google Search” button once they had read and understood the description and query. This was important to ensure that the users started scanning the search results page from the same place as they would if they were entering a query into Google as normal.

As mentioned above, all users saw the same (cached) results page at first. Once they were on this first results page, it was up to them to do whatever they thought they needed to in order to complete the task – e.g. reading text on the results page itself, clicking on links, or changing the query. They could move on to the next task as soon as they felt they were done, or were ready to give up. They did this by pressing the Home button in the web browser, which returned them to their task list.

Before starting on the tasks, each participant was walked through an example and then did a practice task. Each session lasted less than an hour (typically around 30 minutes for all 16 tasks).

Table 1: 4 of the 16 tasks used in the study.

Query	Description
actress most oscars	You are a movie fan and are curious to know which actress has won the most Oscars.
lawn chair	You are going to an outdoor concert soon, and want to get yourself a lawn chair to take with you. Find one online that you would consider buying.
eschew definition	You read an article about healthy eating that listed some foods to “eschew”. You want to check what that word means.
bono real name	You were watching TV and saw something about Bono, the singer. You're curious to know what his real name is.

3. RESULTS

The primary unit of analysis is a *visit* to a Google search results page. Pressing the “Google Search” button to begin a task starts the first visit of the first query. Clicking through on a result and clicking on the Back button in the web browser begins the second visit of the first query (because the results are still the same). In total there were 1216 visits to 786 queries across the 32x16 = 512 tasks. The results reported here apply to all visits in the study.

We used Clearview’s default of 100ms minimum duration and 30 pixel maximum dispersion to determine eye fixations, and used

the raw mouse data for our analyses. We wrote a program that automatically identified the outlines of interesting regions on each results page visited. These were the 10 search results, any sponsored links or additional pieces of information appearing with the results, and the top area of the page (including the search box, logo, etc). The remaining areas of the page were collectively treated as a single region, called “other”.

3.1 Overall Eye-Mouse Relationship

3.1.1 Relative Distribution of Attention

Figure 1 shows the relative distributions of the user’s attention across selected regions of the page, comparing the proportion of mouse data points in each region to the equivalents for total eye fixation duration and total number of clicks. For results 1-10, the distributions are quite similar, but it is interesting to see that the mouse spends a much higher percentage of time in the “other” regions (empty space and the bottom of the page) than the eye.

Using the full set of regions, we were interested to know how likely it was, within a single visit, that when the user moved their mouse over a region, they also looked at it. Of regions that they covered with the mouse, a mean of 76.2% (s.d. 23.4) were also fixated on during the visit. Conversely, of the regions that the users fixated on during a single visit, a mean of 64.0% of those regions (s.d. 25.7) were also covered by the mouse.

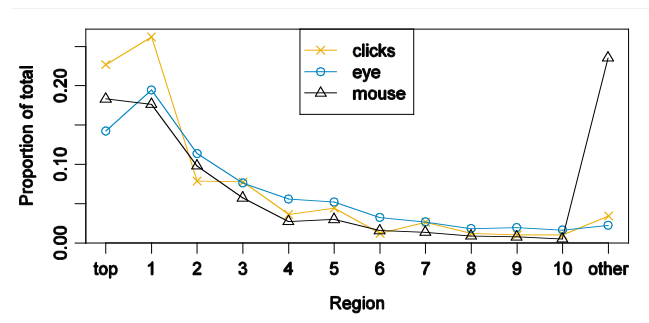


Figure 1: How distributions of eye fixation time and clickthrough relate to distribution of mouse hovering time, for the regions that were common to all pages.

3.1.2 Distance Between Mouse and Eye

In order to calculate the set of distances between mouse and eye, we matched each mouse point with the eye fixation (if any) whose duration spanned it. The overall distribution of distances is skewed, with a long tail to the right. The mean across all fixations is 257 pixels (s.d. 237); the median is 191.

It is interesting to consider the X and Y directions separately (Figure 2). For the Y direction there is a much higher peak around 0 than there is for the X direction, suggesting that mouse and eye positions corresponded more closely in the vertical direction than in the horizontal direction.

Considering the different regions on the page, we found that when the mouse was over results 1-10, the mean eye-mouse distance dropped to 194 pixels (s.d. 132). When it was over the “other” region, the mean rose to 551 pixels (s.d. 305). We have already seen that the participants were more likely to have their mouse in this area than they were to fixate on it. Combined with the data in Figure 2, this might suggest that a common behaviour is to keep the mouse in the blank area to the right of the search results, moving it downwards while scanning the results with the eyes.

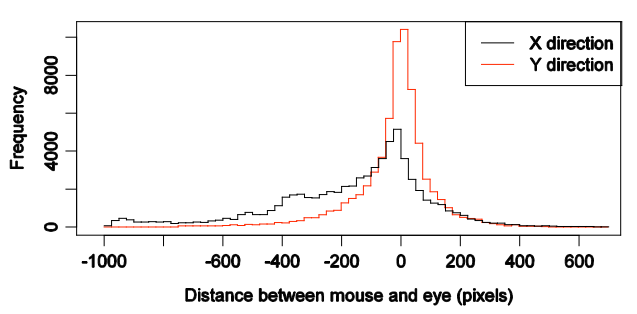


Figure 2: Histogram outlines of distance from mouse to eye, broken out separately for the X and Y directions. Each step in the histogram represents a bucket of 25 pixels. A negative X distance means that the eye was to the left of the mouse, and a negative Y distance means that the eye was below the mouse.

Further evidence for this comes from investigating how frequently the eye and mouse were in the same region at the same time. Overall, eye and mouse were in the same region for 42.2% of the mouse data points that had a corresponding eye fixation and this dropped to 6.0% if the mouse was in the “other” region. So although there is a reasonably high overlap between the regions covered by eye and mouse within a single visit to a results page, this overlap does not necessarily occur at the same time. One exception is the top of the page, where eye and mouse coincide about 70% of the time (when users were refining queries) – this is also the region with the shortest eye-mouse distances.

3.2 Eye-Mouse Coordination Patterns

Following the high-level analysis described in the previous section, we wanted to consider in more detail the interactions between eye movements and mouse movements within a visit. For each visit, we generated a visualization of the paths followed by the user’s eye and mouse. This was straightforward for the eye fixations – we simply placed a circle at the point of each fixation, with area proportional to the fixation duration. To make equivalent visualizations for the mouse data, we created mouse “fixations”, using the I-DT dispersion-based algorithm [12] with a minimum duration of 0ms (so no points were thrown away) and maximum dispersion of 10 pixels. We studied a sample of the visits (see <http://www.rodgen.org/kerry/publications/wisi07/> for example visualizations) to look for instances of the different patterns. In addition to these, we used time series plots to help us understand the relative timings of the events in more detail.

There are some patterns of eye-mouse coordination that we will not discuss further here – in general, these are movements that the user must make in order to complete their task. For example, moving the eye and mouse together to the search box in order to refine a query, or to the scroll bar in order to move further down the page.

3.2.1 Keeping the Mouse Still While Reading

In this pattern, the user holds the mouse away from the place where they are currently reading, keeping it mostly still until they have seen the result they want to click on. In general, the most common starting position of the mouse was the position of the “Google Search” button on the previous page – at the beginning of a task, users were forced to click on this button. This is also the most common starting point for the eye, but the mouse typically

stays in this position for longer than the eye does. So it seems that users very often exhibit this pattern at the beginning of a visit.

As well as the starting position, the blank area to the right of the results (in the “other” region) was also a common resting place for the mouse. Users varied a lot in this regard – the mean percentage of time that they left the mouse in the “other” region ranged from 2% to 57%. Those participants who used the scroll wheel on the mouse would often scroll the page while resting the mouse in this area, resulting in a pattern of evenly-spaced mouse “fixations” in a vertical line. We also saw several pieces of evidence in the high-level data that users adopt behaviours like these at least some of the time, including the mouse being left in the “other” region much longer (relatively) than the eye spends there, and the eye and mouse being at their most distant on average when the mouse is over the “other” region.

3.2.2 Using the Mouse as a Reading Aid

In this pattern, the mouse pointer is moved around to help the user keep their place on the page while reading.

The most common form of this pattern was for the user to move the mouse pointer mostly in the vertical direction, so that it was either touching or roughly level with the region they were currently reading – for the search results, this was often in the “other” area, to the right. This form may help to explain the fact that the mouse and eye tended to be closer in the Y direction than in the X direction.

In another form of this pattern, the user moves the mouse pointer horizontally across or below the text they are currently reading. This is illustrated in Figure 3, and is characterized by sequences “fixations”) while over a result. Although striking, this behaviour was rare in the study. Only a handful of the users ever followed a whole line of text with the mouse while reading it, and they did not do this for every task. It was more likely, however, that in cases where the answer to the task was visible on the page, users would move their mouse over the answer.




Figure 3: Example of “reading” with the mouse (user 21). The user ran the mouse pointer over part of the snippet, containing the answer to the task (“The actress with the most Academy Awards for ‘Best Actress’ is Katharine Hepburn”).

3.2.3 Using the Mouse to Mark an Interesting Result

In this pattern, the user leaves the mouse pointer on or near the result that seems to be the most promising one they have read so far, while their eyes continue to check more results. Often, the mouse is left hovering over the title of the promising result – ready to click if the user eventually decides to select it. If another result seems more promising, the user will move the mouse on to that result, and so on. The difference between this pattern and the previous two is that the mouse is kept still for the purpose of marking an interesting result, not simply to keep it out of the way or to keep the user’s place on the page while reading.

From inspection of the visualizations, and based on previous work on selection from menus [6], we speculate that users will be more likely to exhibit this pattern when the task is more difficult, and it

is not obvious to them which result is best – especially as they move further down the page. Further studies would be required to confirm this.

4. DISCUSSION

The study has given us some tentative answers to the questions we raised in the Introduction. We have found that mouse movements definitely show potential as a way to estimate which results page elements the user has considered before deciding where to click, e.g. by noting which regions were covered by the mouse during the visit, or measuring the vertical distance traversed. This has implications for evaluating the user-perceived quality of the search results (as judged from their surrogates). Mouse movements also have some potential as a method of determining whether the user has noticed the answer to their question on the results page itself, thus helping to evaluate design choices in page formatting and layout. Behaviour such as that illustrated in Figure 3 (following a line of text with the mouse while reading it), is relatively infrequent, but when it does occur, it indicates which aspects of the surrogate the user is taking into account when making a decision.

It is interesting to consider whether it would be possible to automatically identify useful patterns from mouse data alone. For example, the pattern discussed in Section 3.2.3 (using the mouse to mark an interesting result) would be particularly useful to identify, since it indicates which surrogates the user has found most relevant. However, from the mouse data alone, it is difficult to tell the difference between this pattern and that of simply moving the mouse vertically while reading. In both, the mouse pointer is still or relatively still while touching or near a result. Without the eye data, we cannot determine the exact sequence of events: did the user move their mouse to the result simply because they were in the process of reading it, or because they had already read the surrogate and decided it was relevant?

To attempt to narrow this down, we looked at cases where the user was holding the mouse pointer over the title of the result (ready to click), not just on or near the result block in general. We found that in 172 of the 1216 visits (14.1%), the user held the mouse for more than 1 second over the title of a result that they did not click on during the visit. All but one user was represented in these visits. However, manual inspection of a sample of these visits (and the associated visualizations and plots) showed that this heuristic was not enough to separate the two patterns.

In general, automated analysis of this data is complex – even with the most clear-cut examples of the patterns, it would not be straightforward to identify them automatically. Similarly, we found that the users were not easy to classify, and each one seemed to exhibit all of the patterns, to varying degrees. There is also substantial variation between users in all of the high-level measures described in Section 3.1. For example, per user, the mean distance between eye and mouse ranged from 144 to 456 pixels, and the proportion of mouse data points at which the eye and mouse were in the same region ranged from 25.8% to 59.9%.

5. FUTURE WORK

This was an exploratory study, and there is a lot of scope for more research in this area. At a minimum, additional exploratory studies could employ different sets of users and tasks (perhaps having users do their own tasks instead of prescribed ones), as well as different search results page designs.

Controlled experiments, perhaps systematically manipulating the results or the result order according to relevance, would help to confirm some of the findings presented here. In particular, it would be valuable to study the relationship between the information provided by a surrogate and the pattern of marking results with the mouse, e.g. to confirm that if the user keeps the mouse pointer over a particular result while continuing to read others, this is indeed because they saw something relevant in the surrogate.

Findings from such experiments would assist in generating reliable and valid metrics from mouse data. Such metrics would be a prerequisite to conducting any larger-scale studies (e.g. with remote users) where only mouse data can be collected.

6. ACKNOWLEDGMENTS

Xin Fu was an intern at Google when this study was conducted. We are grateful to Chris Mysen, Dan Pupius, Gill Ward, Simon Quellen Field, Laura Granka, Dave Poole, and Tiffany Griffith for their help with this work.

7. REFERENCES

- [1] Agichtein, E., Brill, E., Dumais, S., and Ragno, R. 2006. Learning user interaction models for predicting web search result preferences. In *Proceedings of SIGIR'06*. 3-10.
- [2] Arroyo, E., Selker, T., and Wei, W. 2006. Usability tool for analysis of web designs using mouse tracks. In *CHI '06 Extended Abstracts*. 484-489.
- [3] Atterer, R., Wnuk, M., and Schmidt, A. 2006. Knowing the user's every move: user activity tracking for website usability evaluation and implicit interaction. In *Proceedings of WWW'06*. 203-212.
- [4] Chen, M.-C., Anderson, J. R., and Sohn, M.-H. 2001. What can a mouse cursor tell us more?: correlation of eye/mouse movements on web browsing. In *CHI '01 Extended Abstracts*. 281-282.
- [5] Claypool, M., Le, P., Wased, M., and Brown, D. 2001. Implicit interest indicators. In *Proceedings of IUI'01*. 33-40.
- [6] Cox, A.L. and Silva, M.M. 2006. The role of mouse movements in interactive search. In *Proceedings of the 28th Annual Meeting of the Cognitive Science Society*. 1156-1161.
- [7] Hijikata, Y. 2004. Implicit user profiling for on demand relevance feedback. In *Proceedings of IUI '04*. 198-205.
- [8] Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. 2005. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of SIGIR'05*. 154-161.
- [9] Lorigo, L., Pan, B., Hembrooke, H., Joachims, T., Granka, L., and Gay, G. 2006. The influence of task and gender on search and evaluation behavior using Google. *Information Processing and Management*, 42(4), 1123-1131.
- [10] Mueller, F. and Lockerd, A. 2001. Cheese: tracking mouse movement activity on websites, a tool for user modeling. In *CHI '01 Extended Abstracts*. 279-280.
- [11] Rose, D. E. and Levinson, D. 2004. Understanding user goals in web search. In *Proceedings of WWW'04*. 13-19.
- [12] Salvucci, D. D. and Goldberg, J. H. 2000. Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of ETRA'00*. 71-78.