Deciphering Trends In Mobile Search

Maryam Kamvar and Shumeet Baluja
Google

Understanding the needs of mobile search will help improve the user experience and increase the service's usage. An analysis of search data from a large US carrier showed that cell-phone subscribers are typing longer queries in less time and clicking on more results.

Just as computer-based Web search has been a gateway to increased data consumption, mobile search will help meet the growing user demands for anytime, anywhere data access. With 76 percent of the US population, or 233 million people, subscribing to cell-phone service in 2006 (http://ctia.org/media/industry_info/index.cfm/AID/10323), the potential impact of wireless applications is enormous. Understanding the unique needs of mobile search will help improve the user experience and increase the service's usage.

OVERVIEW OF MOBILE SEARCH

We analyzed data from more than1 million page-view requests randomly sampled from Google logs during a one-month period earlier this year. The requests were anonymous; we maintained no identifying information that could associate searches with users. To eliminate confounding factors between different carriers, we restricted our examination to a single US carrier. To differentiate among computers, PDAs, and cell phones, we looked at the browser's user agent sent in the HTTP request. Unless otherwise noted, the mobile statistics we present pertain to cell phones.

At the time of our study, the Google mobile interface presented users with the option of searching four information repositories: Web (standard Web searches), local (information related to particular geographies), image (keyword-based picture searches), and mobile Web (searches of sites tailored for presentation on mobile phones). To allow accurate comparisons with wired searches, we concentrated our study on Web queries.

We grouped the requests into sessions, which we defined as "a series of queries by a single user made within a small range of time."[1] We referred to this time range as the session delta and used a session time-out of 5 minutes—we deemed a user's session closed if no interaction happened within 5 minutes and considered the next interaction to be the start of a separate session.

A typical search session from a mobile or wired device consists of

- formulating and entering the query,
- browsing the provided search results, and
- viewing the selected result.

Figure 1 illustrates these three steps.

Mobile queries

The average mobile query was 2.56 words (median, 2; maximum, 39; standard deviation, 1.7) and 16.8 characters (median, 15; maximum, 224; standard deviation, 9.2). Interestingly, this was similar to the statistics published for PDA and computer-based queries, where the average number of words per computer-based query reported was 2.35[1,2] and 2.6,[3] and per PDA query was 2.64 (median, 2; maximum, 29; standard deviation, 1.57).

Despite the drastically different input techniques used, the similarity in median and mean query terms across search mediums might suggest that the number of terms

Published by the IEEE Computer Society

per query is currently a *ground truth* for today's Web search. In fact, a small study done on a speech interface to search[4] also found that the average length of spoken queries to Google was 2.1 terms. Users might have learned how to form queries for today's search engines to get neither too many nor too few search results.

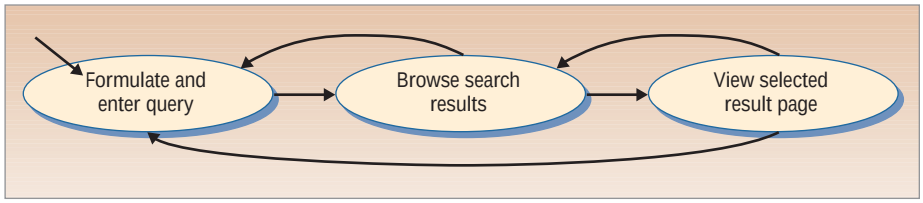

Figure 1. Search-session process. A state diagram illustrates the three steps in a search session.

It's surprising that mobile users don't enter shorter queries given the difficulty of query input. Mobile users have the challenge of entering the query on miniature keypads, most often consisting of a nine-key layout, rather than the conventional qwerty layout. Assuming that users input their mobile queries using the multitap technique, they need an average of 40.9 key presses per query (median, 36; maximum, 720; standard deviation, 1.8).

With multitap, users access letters by repeatedly pressing the key and the system cycles through the letters in the order they're printed. Pausing for a set period of time will automatically select the current letter in the cycle, as will pressing a different key. The amount of effort (the number of key presses) required to enter a word on a cell-phone keypad is more than double the effort required to enter a query on a standard qwerty keyboard.

It takes users a significant amount of time—an estimated 39.8 seconds—to enter these queries. To compute this number, we examined the amount of time between when a user first requests the Google homepage and when Google receives the query request, as Figure 2 illustrates. This number encompasses the time to download the google.com page, input the query, and upload the HTTP request to the server.

The average difference between the two requests (including upload and download time) was about 44.8 seconds (median, 34; standard deviation, 37.8). We subtracted from that 5 seconds to estimate the network latency (the upload and download time) in order to determine the time it took a user to enter a query. Going forward, for all of our estimates of the time it took a user to perform an action, we've subtracted 5 seconds to account for network latency.

As Figure 3 shows, we found that the time to query was proportional to the query length. Furthermore, we found that time to query was also proportional to ease of input. Although queries from PDAs (which often have qwerty keyboards) were longer than queries from cell phones, the average time to input a query decreased to 30.1 seconds.

Since users were willing to spend almost 40 seconds typing their query, the next analysis examined the topics they were willing to spend so much time querying. Table 1 lists the five most popular query categories. The most popular was the adult category, which typically consists of pornographic queries. In comparison, pornographic queries on

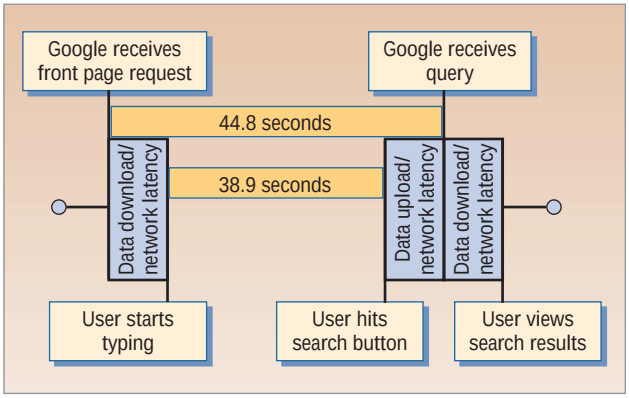

Figure 2. Timeline for query. To account for latency, we subtracted 5 seconds from the time it took to return results.

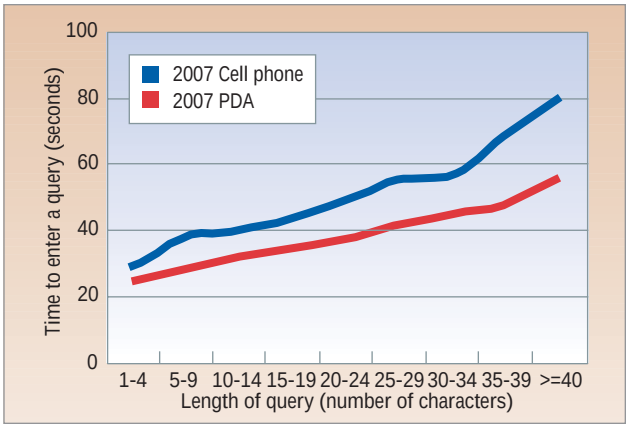

Figure 3. Time to query. Graph of the time it takes to enter a query versus the length of the query.

conventional Web searches accounted for less than 10 percent of all queries in 2001.[3] The same study reported that the proportion of pornographic queries in conventional Web searches declined 50 percent from 1997 to 2001.

We have two hypotheses surrounding the relatively high percentage of pornographic queries submitted in wireless search. First, since wireless search is a more recent phenomenon than desktop search, it could be following the same trend as wired searches. The high percentage of pornographic queries may decline as the service attracts more users.

Second, we speculate that people might feel more comfortable querying adult terms on private devices. The screen is smaller, so it's less likely that a passerby

Table 1. The top five categories in mobile search.

Category	Percent of all queries
Adult	>25
Entertainment	>10
Internet/telecommunications	>4
Lifestyles/online communities	>4
Local	>4
Other	>45

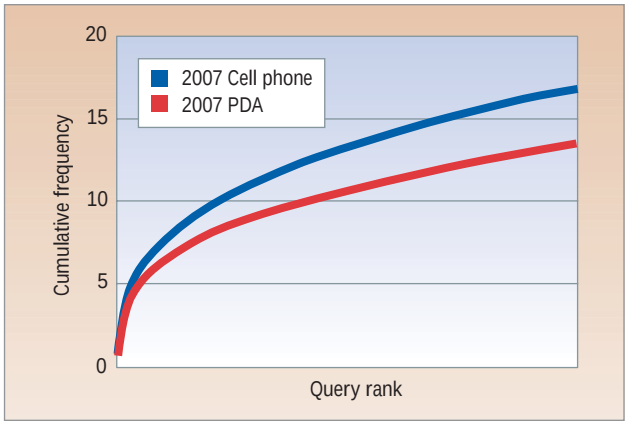

Figure 4. Cumulative frequency of top queries. The frequency of the top 1,000 queries made from cell phones ranked higher than the frequency of the top 1,000 queries made from PDAs.

will notice the nature of the search. Also, users often consider cell phones personal and private, perhaps even more so than their computers. Thus, there's a perceived smaller probability of others discovering their search behavior through cached pages, autocompletion of query terms, or URL history lists.

Examining the distribution of queries across a broad set of topics, as Table 1 shows, is one method to examine the diversity of search requests received. Another measure of the diversity is to examine what percentage of the total query volume the top-*N* unique queries account for. The larger the volume accounted for by the top-*N* unique queries, the less diverse the set of queries received.

To analyze this, we used a random sampling of more than 50,000 queries from cell-phone and PDA searches during a month. Figure 4 illustrates the distribution of the top 1,000 queries. The top mobile query accounted for about 0.8 percent of all wireless queries, and the top 1,000 mobile queries accounted for about 17 percent of all cell-phone-based queries. PDA queries had significantly more variation; the top 1,000 PDA queries accounted for about 13.5 percent of all queries. Computer-based queries are even more diverse. A 2005 study showed that the top 1,000 queries from wired search accounted for only 6 percent of all queries.[5]

One hypothesis for the higher homogeneity of mobile queries is related to the nascent state of the mobile Web itself. People might have adapted their queries to those that return "usable" sites. Usable sites are those that have content that will display well on the search medium (for example, adult content and ring tone sites are "usable" in mobile browsers).

Accordingly, desktop browsers are the most advanced, which would lead to a more diverse set of queries.[5] PDA browsers are less advanced than desktops (they can often display HTML but not JavaScript), and cell-phone browsers are the least advanced, often capable of displaying only limited XHTML content.

A second hypothesis for the decrease in query diversity across wireless mediums is that there's a smaller user base, and that the user base may share similar profiles (for example, cell-phone searchers are likely to be technologically savvy, and PDA users have a business-oriented bent). Following this hypothesis, because desktop browsers are the most readily available and reach the most users, they generate the most diverse queries.

After the query

After issuing a query, the user receives 10 search results. Most users either found what they were looking for on the first page of results or chose not to look further; only 10.4 percent of queries had requests to display more than the initial set of search results.

More than 50 percent of queries led to a click on a search result. It took the average user 30 seconds to scan the search results before selecting one. Of those queries that didn't lead to a click, it's possible that the user found the answer in one of the Web-page summaries returned with each search result, gave up on the search entirely, or refined the search in a subsequent query.

As Figure 1 illustrates, at any point in a search session, a user might choose to modify the original query. The average number of queries per mobile session is 2, (median, 1; maximum, 48; standard deviation, 1.8). Here, we looked at the query pairs that occurred in sessions that had more than one query. Two queries, query 1 and query 2, were considered to be a pair if query 1 occurred before query 2 in the same session. Some 66.3 percent of all query pairs in a session fell in the same category. Furthermore, in all query pairs, the second query was a refinement of the first 58.6 percent of the time. We considered a pair of queries to be a refinement if:

- query 1 was a substring of query 2;
- query 2 was a substring of query 1; or
- the edit distance between query 1 and query 2 was less than half the length of query 2.

From this, we inferred that the majority of wireless searchers approach queries with a specific topic in mind, and their search topics don't often lead to general exploration.

A LOOK BACK

About 18 months have passed between this study and our original study of mobile search in 2005.[5] While that's a short period, we already see a few interesting trends emerging. Table 2 summarizes the statistics.

Table 2. Summary of mobile search statistics in 2005 and 2007.

Mobile search statistics	2005	2007
Words per query	2.3	2.6
Characters per query	15.5	16.8
Percent of queries that had at least one click	<10.0	>50.0
Percent of queries that had at least one "more search results" request	8.5	10.4
Time to enter a query*	56.3	39.8
Time between receiving results and clicking on a spelling correction for a query*	15.6	15.1
Time between receiving results and clicking on a search result*	29.1	30.0

* Assuming 10-second network latency in 2005 and 5-second network latency in 2007

Users type faster

Although mobile queries have slightly increased in length since 2005, the time delta from requesting the Google front page to submitting a query has decreased from 66.3 seconds in 2005 to 44.8 seconds in 2007. We suspect part of this difference is due to shorter network latencies, but we estimated that only 5.5 seconds of the 21.5-second speedup in query entry was due to network improvements.

We estimated the improvement in network latency by comparing the 20.1 seconds it took users to accept a spelling correction in 2007 to the 25.6 seconds it took users to accept a spelling correction in 2005. Since the user interface for spelling correction remained a constant, and since we've noticed that most users are likely to accept a spelling correction without browsing the results, we take the difference in these times to be indicative of the improvement in network latency.

The graphs of query length versus time to enter a query shown in Figure 5 provide evidence that users are typing faster (possibly due to better keyboards or more experience with mobile-phone typing). Note that if network latency were the only factor, we'd expect to see a constant decrease in time to enter a query across query lengths. However, this isn't the case; instead, we observed that the time saved on longer queries was greater than the time saved on shorter queries.

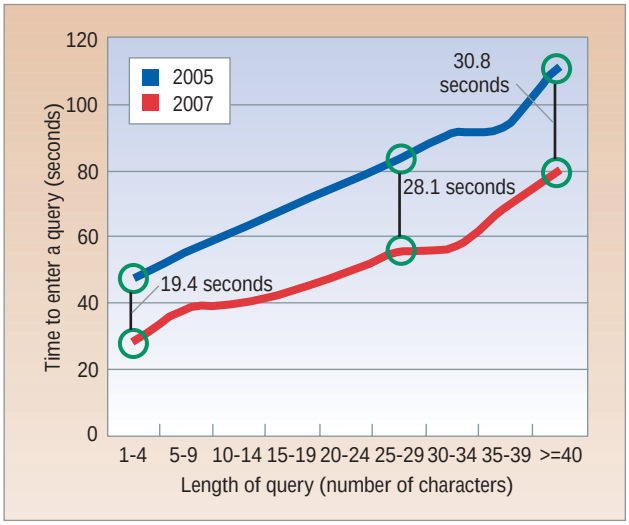

Figure 5. Reduction in query-entry time. Due to faster typing, it took less time in 2007 to enter a query than it did in 2005.

More users are clicking

In 2005, users followed less than 10 percent of queries with at least one click on a search result. In 2007, that percentage rose to well over 50 percent. Additionally, the percentage of queries followed by a request for "more search results" increased from 8.5 percent to 10.5 percent. We attributed the increase in clicks to at least two factors.

First, there have been drastic improvements in the *transcoder* technology that converts a search-results page to a format the user's cell phone can display. In 2005, the transcoder converted HTML to WML, stripping a Web page of all its images and formatting. Now, the transcoder converts HTML to XHTML and retains much more of the formatting and all of the images on the resulting Web page. Second, we believe that the reduction in time to get to the search results (the shorter network latencies and improvement in query-entry speed) have encouraged more users to interact with the search-results page.

Although we find that more users are clicking on the search-results page, the behavior for users who click has remained consistent. The average clicks on search results per query and number of "more-search-result" requests per query are similar in 2005 and 2007.

More exploration within a session

The number of queries per session has increased more than 25 percent from 2005. Although there's low category diversity within a session (most users stick to one category during their search session), we see an increase in query diversity within a session. In 2005, the percent of unrelated consecutive queries was approximately 20-25 percent.[5] Unrelated queries aren't generated by spell-correction suggestions, and they don't classify as query refinements (defined above). In 2007, the number of unrelated queries in a session nearly doubled, to 41.4 percent.

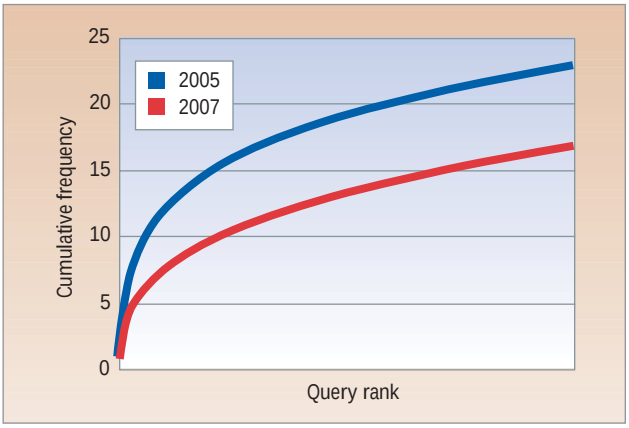

Figure 6. Cumulative frequency of queries. The data compares the frequency of the top 1,000 queries made from cell phones in 2005 to the frequency of the top 1,000 queries made in 2007.

One confounding factor in comparing the two statistics was that in 2005 the measure was taken on consecutive queries, where query 2 occurred directly after query 1 (with no clicks between the two queries). In 2007, the measure was made over query pairs, a less stringent filter where query 2 occurred sometime after query 1 in the same session. However, if we apply the more strict analysis to the 2007 data, we still see an increase: 38.1 percent of consecutive queries aren't related. A partial explanation for this is that the number of identical consecutive queries decreased from 31.7 percent in 2005 to 4.5 percent in 2007.

Less homogeneous queries

As expected, mobile queries are becoming less homogeneous. The top query in 2007 accounted for 0.8 percent of all queries, as opposed to 1.2 percent in 2005. When measuring the cumulative frequency of the top 1,000 queries from a random set of more than 50,000 mobile queries in 2005 and 2007, we observed a decrease from approximately 22 percent to approximately 17 percent, as Figure 6 shows. This may indicate the increasing diversity of mobile Web users and the increased diversity of mobile Web content.

More high-end devices

The percentage of requests from PDAs in the search logs used to account for about 25 percent of the number of requests from cell phones (for the carrier studied). Today, the number of queries from the same carrier originating from PDA devices is about the same as the number of queries from cell phones.

More adult queries

While the relative order and magnitude of query categories remains the same, the percentage of adult queries increased. We attribute the gain to the transcoder improvements, which no longer strips page images after the user clicks on a search-result link. We believe this trend will reverse, as it did with wired queries.

As evidence, we look to the UK, which is often considered more advanced in mobile Web usage. The UK has a much smaller percentage of adult queries. The confounding possibility that UK users are less likely to want adult content is called into question by an anecdotal study which examines the *image-search logs* for both the UK and the US. The percentage of queries related to adult content remains consistent across both countries.

Using anonymous log data, we've presented a brief examination of wireless search patterns for a major US carrier. The strength of such large-scale logs analyses lies in the breadth of data we used. Google is a popular mobile-search site, and analyzing Google's usage provides a wealth of general quantitative information about search traffic. The weaknesses of this method are that these numbers don't tell the story behind a user's experience—we know for what and when a user queried, but have no context for what inspired the search. We also don't know anything about the user's demographics. Despite these caveats, we presented a wide assortment of data on the state of wireless search to provide a useful benchmark in the nascent world of research in this area.

References

1. C. Silverstein et al., "Analysis of a Very Large Web Search Engine Query Log," *SIGIR Forum*, vol. 33, no. 1, 1999, pp. 6-12; www.acm.org/sigs/sigir/forum/F99/Silverstein.pdf.
2. B.J. Jansen et al., "Real Life Information Retrieval: A Study of User Queries on the Web," *SIGIR Forum*, vol. 32, no. 1, 1998, pp. 5-17; http://ist.psu.edu/faculty_pages/jjansen/academic/pubs/sigirforum98/forum98.pdf.
3. A. Spink et al., "From E-Sex to E-Commerce: Web Search Changes," *Computer*, Mar. 2002, pp.107-109.
4. A. Franz and B. Milch, "Searching the Web by Voice," *Proc. Conf. Computational Linguistics* (COLING), Morgan Kaufmann, 2002, pp. 1213-1217.
5. M. Kamvar and S. Baluja, "A Large Scale Study of Wireless Search Behavior: Google Mobile Search," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, ACM Press, 2006, pp. 701-709.

Maryam Kamvar is a software engineer at Google and a PhD candidate at Columbia University. Her research interests include human-computer interaction, small devices, and search interfaces. Contact her at mkamvar@google.com.

Shumeet Baluja is a senior staff research scientist at Google, where he works on machine learning, wireless application development, and user interaction measurement. Baluja received a PhD in computer science from Carnegie Mellon University. Contact him at shumeet@ google.com.