

Good Abandonment in Mobile and PC Internet Search

Jane Li, Scott B. Huffman, and Akihito Tokuda
Google Inc.
janeli,huffman,akihito@google.com

ABSTRACT

Query abandonment by search engine users is generally considered to be a negative signal. In this paper, we explore the concept of *good abandonment*. We define a good abandonment as an abandoned query for which the user's information need was successfully addressed by the search results page, with no need to click on a result or refine the query. We present an analysis of abandoned internet search queries across two modalities (PC and mobile) in three locales. The goal is to approximate the prevalence of good abandonment, and to identify types of information needs that may lead to good abandonment, across different locales and modalities. Our study has three key findings: First, queries potentially indicating good abandonment make up a significant portion of all abandoned queries. Second, the good abandonment rate from mobile search is significantly higher than that from PC search, across all locales tested. Third, classified by type of information need, the major classes of good abandonment vary dramatically by both locale and modality. Our findings imply that it is a mistake to uniformly consider query abandonment as a negative signal. Further, there is a potential opportunity for search engines to drive additional good abandonment, especially for mobile search users, by improving search features and result snippets.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval

General Terms

Measurement, Human Factors

Keywords

good abandonment, mobile internet search, PC internet search, query analysis

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGIR'09, July 19–23, 2009, Boston, Massachusetts, USA.
Copyright 2009 ACM 978-1-60558-483-6/09/07 ...\$5.00.

1. INTRODUCTION

The information retrieval community has a long tradition of using user clicks on search results as a positive signal. Clicks (and sometimes a lack of clicks as well) have been used successfully to learn ranking functions [14, 13, 3, 18] and to evaluate comparative algorithms in A-B [9, 5] or interleaved experiments [8, 10]. It has been considered to be an indicator of user dissatisfaction if users choose not to click on any results, or worse, “abandon” their query by neither clicking a result nor issuing a query refinement [10].

Internet search engines have added features over the past several years that attempt to answer users' information needs directly on the search results page, without requiring a click on any of the results. Leading engines now provide a large array of these features for basic information needs such as weather, stock quotes, local business addresses and phone numbers, images, current news headlines, flight information, package delivery tracking, and many others [1, 2]. In addition, the result snippets returned by search engines have improved over time [19, 17] and may often answer information needs directly [7].

In this paper, we explore the concept of *good abandonment*. We define a good abandonment as an abandoned query for which the user's information need was successfully addressed by the search results page, with no need to click on a result or refine the query.

We present an analysis of abandoned queries sampled from Google's search logs. Specifically, we analyze abandoned queries from three countries (the United States, Japan and China) across two modalities (PC search and mobile search) for a total of six query streams.

We are particularly interested in mobile search and how it compares to PC (desktop/laptop) search with respect to abandonment. We anticipate that there may be differences for several reasons. First, on mobile devices—even current top-tier devices such as the Apple iPhone—opening web pages is often slow and clunky, with formatting issues, usability issues, and content omissions. Therefore we postulate that users might want to avoid opening pages, and instead formulate queries in a way that may return answers directly within search results. Second, anecdotally we hear from users about a “quick answer in a bar” type of use case for mobile search. Here, users are out with friends (and away from their PC), and use mobile search to answer questions that come up in conversation—what's the weather going to be like tomorrow, what time does the movie start tonight, what year was this celebrity born, etc. [16, 12]. This use case, if real, would potentially drive good abandonment on

mobile search. Third, mobile devices are inherently local, and to the extent that internet search engines provide local business addresses and phone numbers, we might anticipate a high rate of good abandonment for queries seeking these types of information.

Our study has three key findings. First, we find that queries potentially indicating good abandonment make up a significant portion of all abandoned queries, ranging from 19% to 55% across the set of locales and modalities we analyzed. Second, we find that the good abandonment rate from mobile search is significantly higher than that from PC search, again across all three locales. This appears to be a meaningful, robust difference between how users interact with mobile search versus search on a PC. Third, by hand-classifying abandoned queries by the type of information need they represent, we identify the major classes of good abandonment, differences across locales and modalities, and perhaps most importantly, largest opportunities for internet search engines to drive additional good abandonment for their users.

The paper is organized as follows. Section 2 summarizes some related work. Section 3 describes the methodology we used to sample and classify abandoned queries, and defines the categories and codings used. Section 4 presents abandoned query analysis results. Section 5 discusses the implications of our findings for internet search engines, and concludes with some pointers to future research.

2. RELATED WORK

Clicks have been treated as the primary implicit user feedback for search engines to learn ranking functions [14, 13, 3, 18] and to evaluate the performance of search algorithms [5, 10]. However, as Joachims, et al. [10] pointed out, more informative than what users clicked on is what they didn't click on. Considering user feedback beyond clicks, "abandonment," describing the user's decision to not to click on any of the results, is of particular interest.

In general, query abandonment has been considered a negative signal, with efforts targeted specifically to reduce it (e.g., [15]). A key contribution of this paper is to introduce the concept of good abandonment. There have been previous studies that imply the existence of good abandonment. For example, Cutrell, et al. [7] conducted an eye-tracking study where web search tasks were broadly divided into two types: navigational tasks and informational tasks, for which results with different snippet lengths were shown to participants. The study suggests that the total time spent on a task was improved for informational tasks when increasing the length of the query-dependent contextual snippet in search results. One implication from this study is that some information needs may have been achieved by viewing snippets without clicking through the results, and therefore leading to good abandonment as defined above.

In this paper, we particularly study the good abandonment across different modalities (PC and mobile) in three locales with significant mobile search usage: the United States, Japan and China. As mobile search is becoming an increasingly important way for internet users to gain access to online information, several efforts have been devoted to analyze the search patterns on mobile devices and compare them to the traditional computer-based search in order to reveal the unique search patterns and ultimately, derive insights for search engines to better serve these users. For

instance, Kamvar, et al. [11] conducted one of the first large scale analysis of English queries from two separate Google mobile logs (Google XHTML and PDA interfaces) and found that users with less sophisticated input capabilities submit shorter queries. This study also highlighted the high percentage of adult-related searches seen in wireless search. Baeza-Yates, et al. [4] conducted a comparison of Yahoo! mobile and computer-based search in Japan, and reported similar query characteristics regarding the query length. A language model based query topic classification in this study suggested the most common query topics in Japanese mobile search were online shopping, sports, and health (in that order), while art, sports, and online shopping were the most common ones for PC search. Another large scale study based on data generated by more than 600,000 European mobile internet users towards the end of 2005 highlighted some important trends in mobile search, such as the shorter average query length, rare use of advanced search features, more limited search vocabulary and higher incidence of repeat queries [6]. A recent logs-based comparison of search patterns across modalities examined the distribution and variability of tasks that users perform, and suggested that search usage is much more focused for the average mobile user than for the average computer-based user [12]. These studies were all large scale analyses based on random query streams, but none focused on abandoned queries.

3. METHODOLOGY AND DEFINITIONS

For the purposes of this study, we define an *abandoned query* as a query that is not followed by any click or any further query within a 24-hour period. We randomly sampled abandoned queries from Google's PC and mobile search logs from a week in September - October, 2008. Following Google's privacy policies, queries containing personal identifying information were excluded from all samples. We also removed mal-formed queries which are nonsense or have encoding errors. We sampled 400 abandoned mobile and 400 abandoned PC queries from Japan (Japanese) and US (English), and 1000 abandoned mobile and 1000 abandoned PC queries from China (Simplified Chinese).

Because our queries are sampled from logs, we don't have access to the user's true information need, or any indication of whether their need was actually met (that is, whether their abandonment was "good"). We therefore have to fall back on our own judgment of potential information needs for each query. We classify a query as a *potential good abandonment* if there is a dominant information need associated with the query that we felt could theoretically be achieved by an internet search engine results page. We made the judgment of *potential good abandonment* for each query by considering the query itself, rather than looking at Google's result page for it. We chose this way in order to obtain an upper bound on good abandonment—counting all queries that could potentially lead to good abandonment if the search results page was ideal, rather than limited by what the search results page looks like today. In other words, this classification attempts to provide an upper bound estimate on good abandonment that internet search engines could generate. In cases where the information need expressed by the query was more ambiguous, but at least one strong interpretation could lead to good abandonment, we simply classify as "maybe" a potential good abandonment. Please see examples in Table 1.

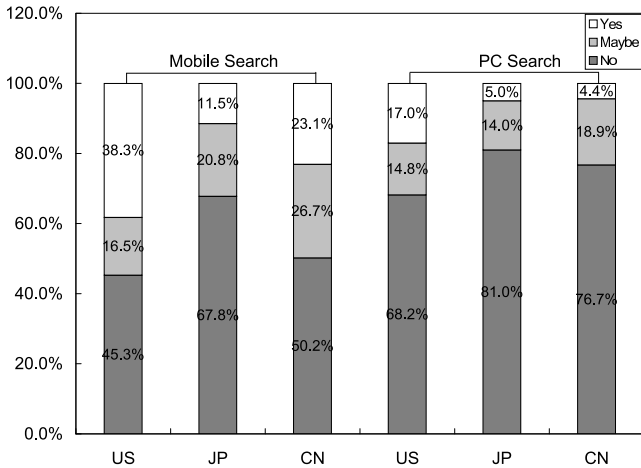


Figure 1: Percentages of queries classified as “Yes”, “Maybe”, “No” with respect to the *potential good abandonment* definition in six abandoned query samples.

We further classify each potential good abandonment query as a *likely* good abandonment by examining the actual results provided by Google. If we felt the information need expressed by the query was clearly met on the results page, we classified it as a likely good abandonment. That is, *likely* good abandonment identifies the subset of potentially good abandonments that were likely to have been actual good abandonments, based on the search results page the user saw. For example, for the query [1 USD in GBP], the results page contained the exchange rates both in a calculator box and in results snippets directly. If we felt the results page may meet the information need but were less sure, we classified the query as “maybe” a likely good abandonment. For example, for [baby come back], partial lyrics were included in result snippets, which may (or may not) satisfy the user’s information need. Table 1 gives more examples.

Finally, the authors coded each potential or likely good abandonment query by category of information need. Rather than pre-defining these categories, each researcher generated their own by doing a first pass through the query samples from their locale. We compared category lists and found that, with slight naming differences, we had converged on a common list. These categories and their definitions are listed in Table 2. Most interestingly, the prevalence and likelihood of good abandonment across the different categories varied significantly both by locale and by modality, as we will describe in detail later in Section 4.3.

4. RESULTS

4.1 Upper Bound Estimate of Good Abandonment

In this section, we report the upper bound estimate of the good abandonment rate by classifying an abandoned query as “Yes”, “Maybe” or “No” with respect to our *potential good abandonment* definition. As shown in Figure 1, “Yes” or “Maybe” potential good abandonments make up a significant portion of all abandoned queries in both mobile and PC searches and in all the three locales we studied.

According to the definition and methodology for identifying a potential good abandonment as described in Section 3, we may estimate the upper bound of the *good abandonment rate* as the *potential good abandonment rate* which refers to the proportion of abandoned queries classified as “Yes” or “Maybe” a potential good abandonment out of all the abandoned queries. For mobile search, abandoned queries whose dominant information need or one strong interpretation could theoretically be answered by a search engine results page count an enormous 54.8% of abandoned queries in the US, 49.8% in China, and 32.3% in Japan. For searches from PC’s, such queries make up 31.8% of abandoned queries in US, 23.3% in China and 19.0% in Japan.

We were surprised by the sheer magnitude of these potential good abandonment rates. Obviously, there are many judgment calls involved in this analysis, and we purposely classified queries with the goal of estimating an upper bound. Even with those caveats, the numbers indicate that good abandonment makes up a very significant portion of abandoned queries.

Comparing query streams, Japanese mobile search presents a significantly lower potential good abandonment rate compared to the other two mobile search locales. Japan is considered to be a more mature market with respect to mobile search [20], perhaps implying larger diversity in search topics and more depth of user’s information needs in this locale. This may be one explanation why we found less good abandonments from the Japanese mobile abandoned queries sample.

As we can also notice from the above statistics, the potential good abandonment rate is significantly higher in mobile search compared to the counterpart in PC search for all the three countries, with a p-value $< 1e-5$ for the large-sample z-test conducted. We theorize that the higher potential good abandonment rate in mobile search may be due to the unique search experience on mobile devices. In particular, 1) Because retrieving web pages over mobile devices is often a clumsy experience, mobile searchers may want to avoid opening pages by querying topics which can be answered directly within search results. We observed a higher rate of simple use cases such as seeking a weather report, a listing of local address and/or phone number, a stock quote, etc., where users can generally expect that the information they seek will be presented directly on the search results page [1, 2]. 2) Mobile information needs are often promoted by contextual factors, which consist of activity at the time, current location and related artifacts, and conversations with others [16]. Many mobile information needs appear to be “quick answer” types of searches, which potentially lead to good abandonment.

4.2 Likely Abandonment Rates

Next, we re-examined each of the queries that were identified as potential good abandonments. We examined the actual results page returned by Google, with the goal of determining whether query’s information need is currently met on the search results page. This gives us a measure of *likely good abandonment rate*—that is, how often the search engine is providing results that likely result in good abandonment for users, as a subset of the queries that potentially could lead to good abandonment. Again, we classified as “Yes” if we felt a query’s information need was clearly met on the re-

Query	Information need	<i>Potential good abandonment?</i>	<i>Likely good abandonment based on Google results?</i>	Comment
[quote MRK]	Stock quote	Yes	Yes	Query was answered by search feature "Stock Quotes".
[weather New York, NY]	Weather report	Yes	Yes	Query was answered by search feature "Weather".
[1 USD in GBP]	Currency exchange	Yes	Yes	Exchange rates were returned in calculator and results snippets.
[taxi buffalo]	Local listings	Yes	Yes	A list of local businesses returned.
[who is the lead singer for tonic?]	A quick answer	Yes	Yes	Answer was provided in title and snippets
[taxi fare nyc from la guardia]	A quick answer	Yes	No	No answer provided on results page
[baby come back]	Lyrics	Maybe	Maybe	Partial lyric was returned in results snippets.
[yamaha psr-172]	Prices/vendors information	Maybe	Maybe	Some price and model information displayed, but may not satisfy the user's need.
[myspace.com]	Homepage	No	No	Search results can provide a link to the desired site, but the "answer" is the site itself
[free ringtones]	Downloads	No	No	The "answer" is presumably actual downloads, which won't surface directly in search results page.
[how to ace an amazon phone interview]	Detailed information on a topic	No	No	The "answer" involves reading more detailed information than would be reasonably surfaced in results snippets.

Table 1: "Potential good abandonment" and "likely good abandonment" query examples

Category	Definition	Query Example
Answer	User seeks a short answer to a question.	[age of consent in PA]
Currency	User seeks currency conversion.	[1 USD in GBP]
Click-to-call	User typed a phone number into search, and either meant to call it, or seeks to learn whose number it is.	
Definition	User seeks the definition of a term.	[what is nvmd]
Images	User wants pictures of a person or thing.	[cubs logo]
Local	User seeks a local listing (address and/or phone number).	[at&t wireless, bartlett, tn]
Lyrics	User seeks song lyrics.	[abettes with time lyrics]
Map	User seeks for a map of a location (address or geographic locations).	[E 83rd St, Los Angeles, CA 90001]
Celebrities	User seeks news or images of a celebrity.	[john stamos]
News	User seeks current news on a topic.	[santa cruz wild fires]
Product	User seeks simple product information such as price range and typical vendors.	[2000 gsx 750f for sale]
Person	User is looking for contact information or vanity information on a (non-celebrity) person.	
Quotation	User seeks for the reference of a quoted sentence/phrase.	[“(a) %20 aqueous phosphoric acid (H3PO4) or (b) an aqueous solution]
Stock	User seeks a current stock price.	[quote AKAM]
Sports	User seeks sports scores.	[chicago cubs]
Showtimes	User seeks movie showtimes.	[The Dark Knight]
Spelling	User seeks correct spelling of a word.	[unfortuanatly]
SMS	User seeks short (greeting) messages to send.	[Short greeting message] (in Chinese)
Translation	User seeks for the translation of a word in foreign language.	[dounika translation]
Weather	User seeks a weather report.	[Weather New York]

Table 2: Task categories and their definitions




Figure 2: Percentage of potential good abandonment queries which are classified as Yes/Maybe/No with respect to the likely good abandonment definition

sults page, “Maybe” if we were less sure or there was partial information, and “No” otherwise.

Figure 2 shows the results. The full bars represent the percentage of abandoned queries that are potential good abandonments in each query stream. The bars are divided into Yes/Maybe/No categories of likely good abandonment. That is, the lightest grey bars indicate the portion of abandoned queries that we felt were likely to be actual good abandonments; the darker grey bars may have been, but we were less sure based on the information shown on the results page. The black bars are perhaps the most interesting. They indicate clear headroom—queries that have the potential to be good abandonments, but whose information needs are not addressed on the search results page today.

Looking at this graph, we can see again that the rate of potential good abandonment is significantly higher for the mobile query streams than the PC query streams, in all three locales. However, the headroom for driving additional good abandonments (the black bars) is much closer across PC and mobile, hovering between 8% to 15% of abandoned queries for all six query streams.

In other words, the search engine is successfully “answering” a greater proportion of the mobile queries that are potentially answerable on the results page. For PC search, an average of 56% of potential good abandonments were clearly or possibly met on the results page. For mobile search, that number is 70%.

This doesn’t imply that mobile search results have higher quality. Rather, it may be attributed to mobile searchers having more focused, less complex information needs than PC searchers. Other studies on random mobile query streams indicate this (e.g., [12, 6]), and we noticed similar patterns in the abandoned query streams we analyzed. For example, we observed that 18% of potential good abandonments in Chinese mobile search were weather queries (a simple information need), while on Chinese PC search the rate was under 1%. Similarly, 43% of US mobile search potential good abandonments appeared to seek local addresses or phone numbers, compared with under 28% of potential good abandonments from US PC search.

Comparing locales, Chinese mobile search has the lowest headroom for driving additional good abandonment (just over 9% of abandoned queries). Of potential good abandonment queries, in Chinese mobile search over 80% are fully or partially answered on the results page. Again, this appears to be because mobile queries in China target relatively simple information needs that are addressed directly by existing search engine features. As the mobile search market grows and mobile devices become more capable, this may change.

4.3 Classification by Information Need

In this section, we further classify potential good abandonment queries by the type of information need they express. The goal is to identify the major information needs expressed by queries which potentially lead to good abandonment, and to discover which categories have headroom to drive additional good abandonment. For the rest of this section, we will restrict our analysis to potential abandonment queries.

As described above, we allowed categories to emerge as we coded queries, rather than predefining them, and converged on the list of categories shown in Table 2. All queries labeled as “Yes” or “Maybe” a potential good abandonment were coded. The percentages of each category of the six potential good abandonment query streams are shown in Figure 3.

As indicated by Figure 3, the distribution of good abandonment information needs varies enormously across search modalities and locales. Queries seeking local information or short answers are the top classes leading to good abandonment in PC search, consistently across locales. There are significant portion of entertainment related searches in Japan and China, especially in mobile search.

The large differences across query streams can be seen by listing categories covering over 10% of potential good abandonment. For US mobile search, they are Local (42.9%), Answer (22.3%), and Stock (11.9%); for Japanese mobile search, Local (24.8%), Answer (20.1%), Celebrities (17.8%) and Images (13.9%); and for Chinese mobile search, Weather (17.9%), Answer (14.7%), Celebrities (14.5%) and News (14.3%). For PC search, the major classes in US search are Local (28.4%) and Answer (20.5%); in Japan, Celebrities (25%), Definition (21.1%), Local (18.4%) and Answer (15.8%); and in China, Local (25.8%) and Answer (22.3%).

The different category distributions across modalities and locales reveal some interesting patterns. Figure 4 highlights five categories worth exploring further.

“Local” search appears to be a large good-abandonment class, with high potential for driving additional good abandonments in all of the query streams except Chinese mobile. Example (US) queries include [domino’s pizza -cary, nc] and [kohls in austin]. One surprising finding is that there are less than 3% “Local” potential good abandonment queries for Chinese mobile search. By contrast, it is the top class for Chinese PC search. By examining the actual search results of these Chinese “Local” queries, we found that they were poorly addressed on the search results page in many cases. It may be that we see fewer abandoned “Local” searches on Chinese mobile search because these searches often do not return the specific information asked for on the results page. Mobile users may have less tolerance because of the extra effort involved to search on mobile devices, where text entry is difficult and the time of loading a page gets longer. Perhaps mobile searchers “give up” more quickly on search-

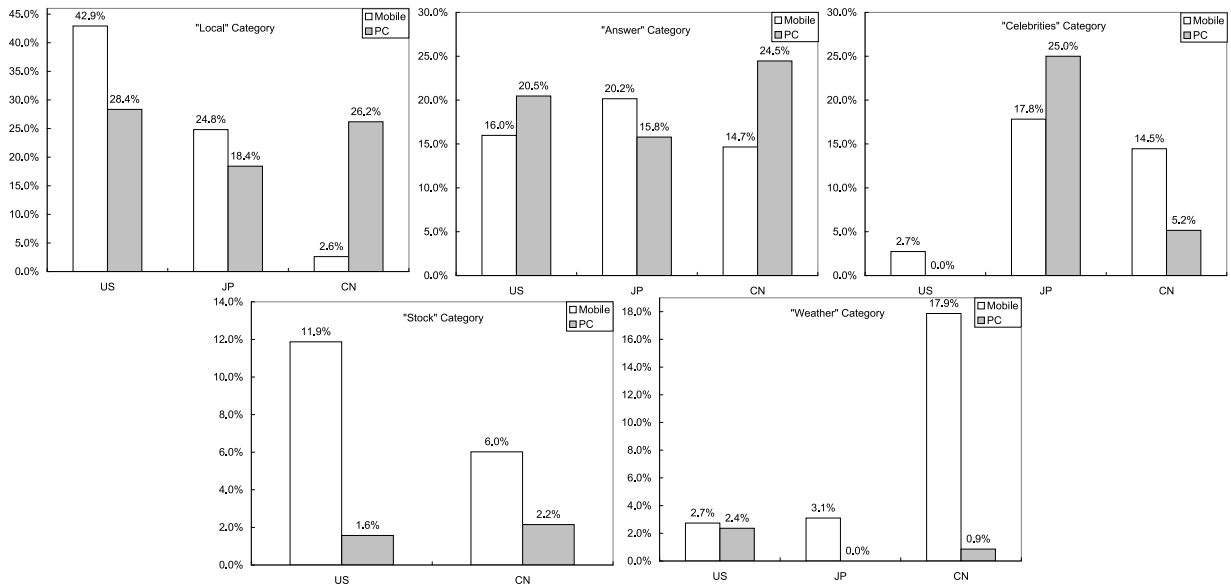


Figure 4: Categories with varying prevalence across modalities and locales.

ing categories for which they previously had an unsuccessful experience.

“**Answer**” is a major class of good abandonment queries across all six query streams. In reviewing the search result pages for these queries, we found that both result snippets and search features such as calculator, definition and news headlines are rich sources for addressing this type of information need. As search engines improve in these areas (introducing more features and providing more intelligent snippets), more “Answer” queries can potentially lead to good abandonment. This has implications for how these types of improvements are evaluated—driving more clicks is not always a success.

“**Celebrities**”, where a user seeks for news and/or images of famous people, is most prevalent in Japanese abandoned queries, but rare in the US. According to several previous studies of random query streams, celebrity searches, which might fall into “entertainment” or “news and current event” categories in these studies, are popular in US [11, 12] and Japan [4] for both PC and mobile searches.¹ The uneven distribution of “Celebrities” search in our abandoned query streams therefore indicates different user interaction patterns by locale and modality for celebrity topics. Presumably, US users click more often when searching celebrities, while Japanese mobile users apparently often browse images or news outlines appearing on the results page.

“**Stock**” is another interesting category which may indicate some differences between the two modalities. While not appearing at all in our Japanese abandonment samples, in the US and China, its prevalence is consistently higher in mobile search. This difference may stem from two factors: First, PC search in general involves larger diversity and less focused search usage. Secondly, the depth of user’s information needs may vary across modalities. For example, a mobile user may just browse the stock quote returned on the results page, since loading additional pages to explore further is slow on mobile devices. In contrast, a PC user may

click to seek for more information and graphs, and therefore a lower fraction of searches are abandoned.

Finally, “**Weather**” search is prevalent only in China mobile search. Weather searches account for 18% of Chinese mobile potential good abandonments, but less than 1% of Chinese PC good abandonments, and less than 3% of mobile and PC in the other two countries. The large count of simple weather queries may correlate with China being a relatively less mature mobile search market compared to the other two locales.

Next, we turn our attention to exploring which categories present the most opportunities for driving additional good abandonment. By counting queries that are potentially good abandonments, but are not fully answered by the current search results page, we can get a measure of the “headroom” for driving further good abandonment in each information need category. Figure 5 shows a heat map of this headroom by category in each locale and modality. Each column represents a locale and modality, and each row represents a category. Each cell represents a number from 0 to 100, which is a normalized count of “No” and “Maybe” likely good abandonments in the category. In our count, each “No” query counted as 1.0, and each “Maybe” as 0.5. In the heat map, the darker cells represent a count closer to 100 (more headroom), and the lighter cells indicate the counts closer to 0 (less headroom).

As suggested in Figure 5, the current search results page serves fairly well for relatively simple information needs such as “weather”, “stock”, “showtimes”, and “spelling”. Several other categories show room for improvement, such as the “Local”, “Answer”, “Images” and “Celebrities” categories. There is an opportunity for search engines to directly address “Local”, “Images”, and perhaps “Celebrities” information needs more often on the results page through more aggressive triggering of corresponding “onebox” or “shortcut” insertions. For the “Answer” category, more intelligent snippets would benefit users, especially mobile users, who are seeking for a short answer or news update. It may make sense to display longer snippets that provide more topical information

¹We don’t know of similar analyses for Chinese queries.

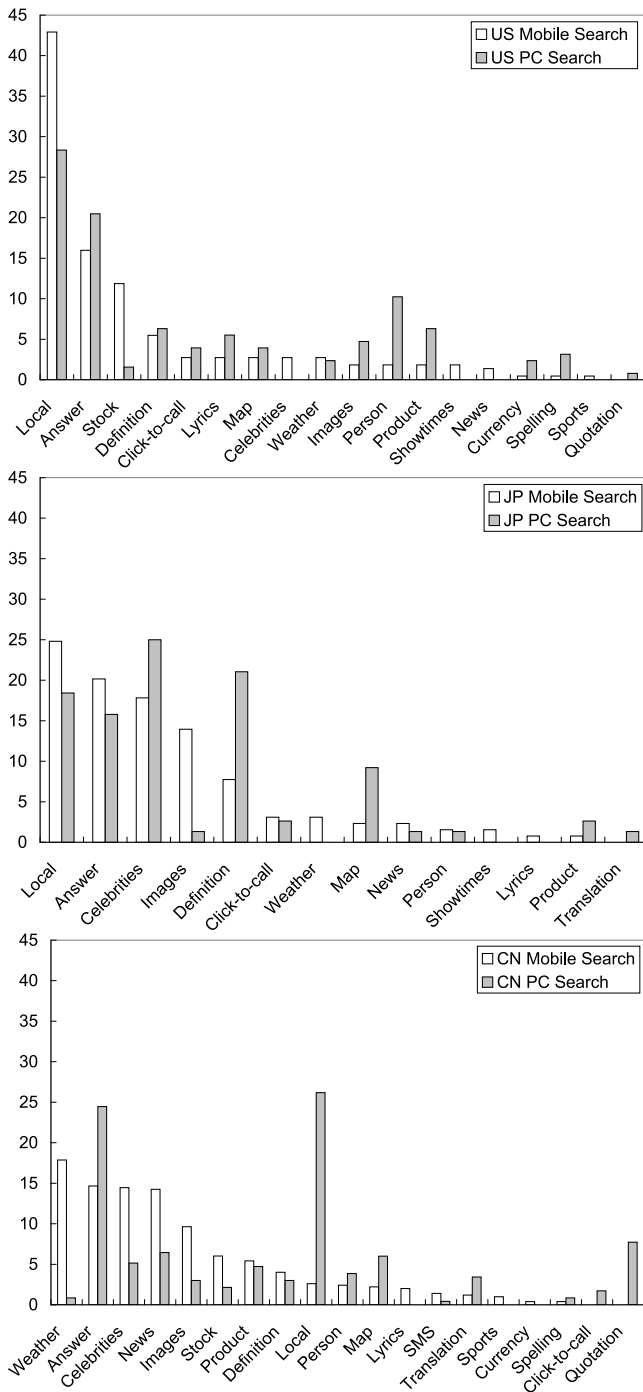


Figure 3: Category distribution (in percentage) of potential good abandonment queries in mobile and PC searches in three countries. The categories are sorted by their prevalence in mobile search for each locale.




Figure 5: Heat map depicting the headroom to drive additional good abandonment by information-need category. Darker squares indicate more headroom.

in mobile search, since loading a web page is often a painful experience on mobile devices.

5. CONCLUSIONS AND FUTURE WORK

Our study has several findings. First, we have shown that queries potentially indicating good abandonment make up a surprisingly large subset of all abandoned queries. Second, across each of the three locales analyzed, we observe a good abandonment rate on mobile search that is significantly higher than PC search. Third, broken down by type of information need, the major classes of good abandonment vary widely by locale and modality. While some classes are large across all query streams, such as the “Local” and “Answers” categories, others appear significantly only for a specific query stream, such as the “Definition” category for Japanese mobile search, and the “Weather” category for Chinese mobile search. Finally, we have identified the categories in each query stream with the most headroom for driving additional good abandonment.

For mobile search, our analysis didn’t differentiate searches from top-tier phones, such as the Apple iPhone, from that of conventional mobile phones. Along the same lines, we didn’t differentiate among data connection speeds, which vary widely by mobile device and carrier network. Considering that there are likely different search behaviors on higher-end devices, as claimed in a recent study [12], it would be worthwhile to examine abandonment patterns separately on these devices in the future study.

We used our judgment to hand-classify queries into potential and likely good abandonment buckets and information need categories. While we consider this a more accurate way to classify queries than the automated topical classi-

fiers used in most larger-scale analyses [4, 11, 6, 12], the approach is clearly limited. It would be ideal to conduct a large-scale diary study with real search users to get a true gauge of good abandonment from the users' perspective, and to understand the true information needs involved.²

One implication of this study is that it is a mistake to uniformly consider query abandonment as a negative signal. Abandonment is no longer always a reflection of user dissatisfaction. Given the surprisingly large rates of potential and likely good abandonment, it may be that future ranking and evaluation models derived from clicks can be improved by taking good abandonment into account.

Another implication is that there is still an opportunity for search engines to directly address users' information needs more often by providing the right information on the search results page. That is, for several types of information needs, there is reasonable headroom to drive additional good abandonment.

This is especially true for mobile search, where there appears to be higher demand for good abandonment in the query stream. Similarly, a recent user study concluded there are more queries per session on computers than on mobile phones, implying that mobile users are less willing to explore topics in depth given the larger barriers to exploration, such as difficult query entry and network latency [12]. For mobile search, perhaps a richer results page with longer or higher quality result snippets, more aggressive triggering of "onebox" or "shortcut" results, or more categories of these types of results, would better serve mobile search users.

6. ACKNOWLEDGMENTS

We would like to thank Neha Arora, Sameer Shariff and Takeshi Yoshino for helping us with data collection, and Dan Russell for helpful comments on an earlier draft.

7. REFERENCES

- [1] Google search features: <http://www.google.com/intl/en/help/features.html>.
- [2] Yahoo! shortcuts services: <http://help.yahoo.com/l/us/yahoo/search/basics/basics-05.html>.
- [3] E. Agichtein, E. Brill, and E. Dumais. Improving web search ranking by incorporating user behavior information. In *Proceedings of the 29th annual international ACM SIGIR*, pages 19–26. ACM, 2006.
- [4] R. Baeza-Yates, G. Dupret, and J. Velasco. A study of mobile search queries in japan. In *Proceedings of the International World Wide Web Conference*, 2007.
- [5] B. Carterette and R. Jones. Evaluating web search engines using clickthrough data. *Submitted to SIGIR*, 2007.
- [6] K. Church, B. Smyth, P. Cotter, and K. Bradley. Mobile information access: A study of emerging search behavior on the mobile internet. *ACM Transactions on the Web*, 1(1), 2007.
- [7] E. Cutrell and Z. Guan. What are you looking for?: an eye-tracking study of information usage in web search. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 407 – 416, 2007.
- [8] T. Joachims. Evaluating retrieval performance using clickthrough data. In *Proceedings of the SIGIR Workshop on Mathematical/Formal Methods in Information Retrieval*, 2002.
- [9] T. Joachims, L. A. Granka, B. Pan, H. Hembrooke, and G. Gay. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th annual international ACM SIGIR conference*, pages 154–161. ACM, 2005.
- [10] T. Joachims and F. Radlinski. Search engines that learn from implicit feedback. *Computer*, 40:34–40, 2007.
- [11] M. Kamvar and S. Baluja. A large scale study of wireless search behavior: Google mobile search. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 701–709, 2006.
- [12] M. Kamvar, M. Kellar, P. R., and Y. Xu. Computers and iphones and mobile phones, oh my! a logs-based comparison of search users on different devices. *WWW*, 2008.
- [13] F. Radlinski and T. Joachims. Query chains: learning to rank from implicit feedback. In *Proceeding of the 11th ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 239–248. ACM, 2005.
- [14] F. Radlinski and T. Joachims. Active exploration for learning rankings from clickthrough data. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 570–579. ACM, 2007.
- [15] A. Sarma, S. Gollapudi, and S. Ieong. Bypass rates: Reducing query abandonment using negative inferences. In *Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 177–185, 2008.
- [16] T. Sohn, K. Li, W. Griswold, and J. Hollan. A diary study of mobile information needs. In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems*, pages 433 – 443. ACM, 2008.
- [17] A. Turpin, Y. Tsegay, D. Hawking, and H. E. Williams. Fast generation of result snippets in web search. In *Proceedings of the 30th annual international ACM SIGIR conference*, pages 127–134. ACM, 2007.
- [18] G. R. Xue, H. J. Zeng, Z. Chen, Y. Yu, W. Y. Ma, W. S. Xi, and W. G. Fan. Optimizing web search using web click-through data. In *Proceedings of the thirteenth ACM international conference on Information and knowledge management*, pages 118–126. ACM, 2004.
- [19] X. Xue, Z. Zhou, and Z. Zhang. Improve web search using image snippets. In *Proceedings of the National Conference on Artificial*. ACM, 2006.
- [20] K. Yamauchi, C. W., and W. D. A study on japanese mobile phone market and its applications. In *Computer and Information Technology, The Fourth International Conference*, pages 875– 878, 2004.

²One of us was recently watching a basketball game in which IUPUI was playing. He pulled out his phone and searched for [iupui]—a navigational query that we would not have classified as a potential good abandonment. The results page met his information need (learning the name of the university) and he abandoned the query. As this example shows, even seemingly clear cases of "bad" abandonment can actually be good abandonment, and vice-versa.