Characterizing Landing Pages in Sponsored Search

Haibin Liu
The Pennsylvania State University
University Park, PA 16802, USA
haibin@psu.edu

Woo-Cheol Kim
The Pennsylvania State University
University Park, PA 16802, USA
wxk11@psu.edu

Dongwon Lee
The Pennsylvania State University
University Park, PA 16802, USA
dongwon@psu.edu

***Abstract*—Using a total of 60,419 ad links collected from three search engines (i.e., Bing, Google, and Yahoo), we characterize the "mobile-friendliness" of landing pages in sponsored search. In particular, we analyze the common and different characteristics between landing pages made for desktop vs. mobile device users, measure/validate the quantitative scores for their mobile-friendliness, and classify the results with respect to types of queries and landing pages. Based on our findings, we articulate that: (1) current landing pages (regardless of search engines or platforms) are *not* mobile-friendly enough, and (2) better data-driven methods (as opposed to current static methods) to help advertisers build mobile-friendly landing pages are needed.**

I. Introduction

When a user clicks an ad that is displayed as part of sponsored links in search engines' advertisement programs (e.g., Google's AdWords and Microsoft's adCenter), she is redirected to a web page called *landing page* (LP) that contains relevant extension of the ad in the advertiser's web site. Since such landing pages often form the first impression about the advertisers (and their products or services thereof), their quality has a major implication on the success of the advertising. The quality of landing pages becomes even more important when it comes to mobile usage. According to Morgan Stanley's forecast in 2009 [1], for instance, mobile usage will be at least double that of the desktop/laptop within the next 5 years. Therefore, when advertisers design their landing pages for online advertising, it becomes increasingly more critical to make landing pages "mobile-friendly." Our focus, in this paper, is more pertinent to *traditional* web browsing on mobile devices, excluding other content presentation options or emerging multimodal technologies on mobile devices (e.g., multimedia messaging, ringtones, podcasts, and mobile apps). First, let us use the following informal definition:

Definition 1 (Mobile-Friendly Landing Page) *A landing page that can be rendered well in mobile devices is called as the mobile-friendly landing page.* □

Note that the definition of mobile-friendliness is vague and subjective at best. Often, whether a page is going to be displayed well in a mobile browser or not depends on many factors. W3C's Mobile Web Best Practices[1] suggests that mobile-friendliness of a page be related with the types of *content* involved (e.g., narrow image vs. flash-based animation), the *capabilities* of mobile devices and networks used (e.g., basic cellular radio access vs. 4G LTE), and the *context* in which the content is received by the user (e.g., sitting at a desk vs. standing in a subway).

Regardless of the precise definition of mobile-friendliness, businesses deal with the issue differently. Some businesses make their landing pages (and web sites) such that they render well across multiple devices. Other businesses maintain separate contents (and URLs to such contents): one for desktop and the other for mobile devices. Conventionally, main web site URLs such as "www.foo.com" are used for desktop contents, while special URLs such as "m.foo.com" or "foo.com/m/" are reserved for mobile contents. To differentiate these two types of landing pages, in this paper, we use the following terms:

Definition 2 (Desktop/Mobile Landing Page) *When a search engine returns a landing page for user agents from fixed devices (e.g., desktop, laptop, tablet), the returned landing page, denoted as LP_d, is called a* ***desktop*** *landing page. Similarly, when a search engine returns a landing page for user agents from hand-held communication devices (e.g., cell phone, smartphone), the returned landing page, denoted as LP_m, is called a* ***mobile*** *landing page.* □

By and large, desktop landing pages are designed for bigger screen sizes than mobile landing pages are. Note that our definitions of desktop vs. mobile landing page are only determined by "platforms" on which sponsored search is made. Therefore, it is possible that an exactly identical landing page may be viewed as either desktop or mobile if returned for, for instance, iPad or iPhone users, respectively. Despite the importance of mobile-friendliness of landing pages, in general, very little is known about their characteristics. Although some preliminary work exist on computational advertisement in general and landing page classification in particular, no existing work studied the landing pages from the mobile-friendliness stand point. In this paper, therefore, we present (to our best knowledge) the first study to identify several important characteristics of both desktop and mobile landing pages, and their "mobile-friendliness" results.

II. Related Work

There exist tools to check mobile-friendliness of a web page according to common practices (e.g., mobiReady, Google Page Speed). However, there is currently no fundamental understanding as to features of landing pages affecting their mobile-friendliness. For instance, the Google Mobilizer, a chrome

[1] http://www.w3.org/TR/mobile-bp-scope/

extension, can instantly generate mobile-friendly version of a web page. While it generates a decent quality output, often, for a page with complex internal structure/graphics, all it does is to strip off textual contents from the page. Therefore, commercial advertisers who want to keep the graphical design aspect of landing pages will not find Google Mobilizer acceptable. Using CSS3 Media Query, a designer is able to extend a single page design across desktop, tablet, and mobile devices. However, it helps little in rewriting existing mobile-unfriendly landing pages into mobile-friendly ones. Commercial tools (e.g., WireNode, Mobify, bMobilized, Onbile) are in abundance to help advertisers create mobile-friendly web sites. However, many of them focus on creating mobile web sites from the scratch (as opposed to rewriting existing ones), or require intensive labor by web designers.

In academic literature, there have been a few attempts to adapt existing web pages for mobile devices (e.g., [5], [7], [13], [14]). Although useful, none of them provides quantitative scores w.r.t. how similar a page is (before and after the rewriting). [4] studied the causal relationship between landing pages types and their conversion rates. They reported that landing pages could be classified into three major distinct types–*Homepage*, *Category browse*, and *Search transfer*–and a supervised classifier with high accuracy can be built. Normally for sponsored search, search engines would always base the ad selection on *bid phrases*, which are specified by advertisers to maximize the response of target audience. As a result, selecting right bid phrases becomes a non-trivial burden for advertisers. Furthermore, often, the contents of landing pages may not sufficiently meet users' information need. To alleviate such problems, [6] proposed different strategies to compute the similarity of feature vectors composed of bid phrases, ad titles and descriptions. [12] used landing pages to automatically generate relevant and well-formed bid phrases. The authors experimented with translation models, content match system, and discriminative system for bid phrase ranking.

Besides computational advertising, mobile web browsing is also relevant. Because of the limitations in mobile handheld devices, including small screen size, narrow network bandwidth, low memory capacity, and limited computing power and resources, researchers explored different methods to improve loading and visualizing large documents on handheld devices. For instance, [8] discussed how to avoid distorting web pages in mobile devices using segmentation of contents and ranking therein. [2] argued that changing the layouts of web pages would simplify or delete contents of the pages, leading to undesired misunderstanding. Instead, the authors designed systems to facilitate users' browsing experience. Recent efforts (e.g., [11], [10], [3]) studied the impact of mobile web in developing countries or for blind people.

III. Design of the Study

A. Collecting Landing Pages

We pick three major search engines that currently support sponsored search–Bing, Google, and Yahoo. All three search engines have separate search URLs for desktop and mobile agents. Using 800 labeled queries in the 2005 KDD Cup data set [9] as encoded URL parameters, we submit HTTP GET requests to six search URLs (of three search engines). We set the values of user agent appropriately to emulate different behaviors for desktop and mobile platforms. Then, we collect all ad links from the *first* page of returned search results. Ad links in sponsored search are usually found on the top and right of a page and sometimes on the bottom. From the collected ad links, next, we extract URLs to landing pages. Note that URLs to landing pages are often embedded within ad links in some hard-coded manners which vary across search engines. After de-duplication and cleaning steps, finally, we obtain a set of landing pages for both desktop and mobile agents. We interchangeably refer to data sets containing landing pages returned for desktop and mobile users as `desktop` (i.e., LP_d) and `mobile` (i.e., LP_m), respectively. Overall, in our experiments, we collected and analyzed a total of 60,419 landing pages across three search engines and two platforms (LP_d and LP_m) from July to November, 2011.

B. Testing Mobile-Friendliness with `mobileOK`

The mobile-friendliness of a given landing page (i.e., whether the page can be rendered well on a mobile device) may change dramatically, depending on the contents, capabilities, and contexts of the evaluation. For instance, the same landing page may be displayed well in the latest iPhone with 4G but not so in a barebone cell phone with poor network bandwidth. Therefore, inherently, it is challenging to test whether a given landing page is mobile-friendly or not. There are several tools to test mobile-friendliness of web pages or sites (e.g., mobiReady, Gomez, iPhoney). In this paper, among these, we decided to use the W3C `mobileOK` checker that enables machine-verifiable interface (as opposed to web interface) and focuses on the mobile-friendliness of "entry-level" mobile devices (instead of feature-rich smartphones). This checker performs around 100 sub-tests based on the `mobileOK` Basic Tests 1.0 specification[2] and the Mobile Web Best Practices[3]. In particular, `mobileOK` tests different areas that impact the mobile-friendliness of page contents such as markup validation, structure/size of the page, CSS style sheets, user input, HTTP level, and links. Although the `mobileOK` basic conformance does not automatically imply an endorsement or suitability of contents for all mobile devices, it primarily assesses the very basic usability, efficiency and interoperability of contents.

In our experiments, we use `mobileOK` checker library v 1.4.2. Once a URL to a landing page is given, `mobileOK` produces an XML document with detailed test results about success or failure of each test and a mobile-friendliness score of 0–100 range. Then, using the scores from `mobileOK`, we define the following:

Definition 3 (MF-score) *MF-score, i.e., Mobile-Friendliness score, of a landing page p refers to a score that* `mobileOK`

[2]http://www.w3.org/TR/mobileOK-basic10-tests/

[3]http://www.w3.org/TR/mobile-bp/

assigns to p (after scaled down to 0–1 range). □

The MF-score returned by `mobileOK` checker is computed based on the number and severity of failures of 100 sub-tests carried out over a web page. Each failure can be diagnosed in a severity level between 1 (low) and 6 (critical). More severe failures will cause more penalty in the MF-score evaluation.

If a landing page gets an MF-score of 1, it implies that the page is likely to be laid out well in a barebone cell phone. Reversely, the MF-score of 0 means that most mobile devices will not be able to render (part of) the page or will not be able to render the page in a reasonable time frame. In addition, the MF-score given by `mobileOK` is not determined linearly. For instance, the MF-score of 0.5 does not imply the passing of half of 25 tests in `mobileOK` nor suggest 100% more mobile-friendly than that of 0.25.

C. Impact of Time and Location

The contents and ad links of sponsored search often vary even for the same query. In general, depending on factors such as personalization on user side, time/location of the query, and specific ad selection/ranking algorithm, the list of returned ad links may differ significantly. For instance, for a query "bike", sponsored search may return a few local bike shops in the New York city at high ranks if the experiment was done in lower Manhattan. However, if done in a small town like State College, PA, the same query may return only national bike manufacturers or retail shops since not many local businesses are likely to participate in online advertisement programs. To study the impacts of such factors on our empirical study, therefore, we set "no personalization" on user side and focused on two factors: *time* and *location* of the query. First, we repeated the same experiment (e.g., extraction of landing pages for the same query) via multiple data points across several days apart, and compared their differences through all subsequent experiments. Second, using Amazon.com's cloud computing environment[4], we concurrently ran the same set of experiments from two physical locations–US-west (North California) and US-east (North Virginia), and compared their results. Third location–State College (Central Pennsylvania)–is also used in the experiments.

At the end, our experiments reveal that even if individual ad links may change depending on time and location of the experiment, the overall patterns (e.g., average # of returned ad links and their average MF-scores across locations) remain the same. Therefore, in the interest of space, we do not present all results with varying time and location in experiments, and instead focus on main findings in subsequent sections.

IV. Experimental Results

A. Setting the Ground Rules

First, we attempt to validate whether or not `mobileOK` is effective in quantifying the mobile-friendliness of a landing page. Our challenge lies on that there has been no benchmark "ground truth" of mobile-friendly and mobile-unfriendly landing pages. Therefore, we propose to simulate the ground truth data set in two ways–one based on *heuristics* and the other based on *human judges*. We based our evaluation on the top-500 most-visited US sites from Alexa[5].

TABLE I
Comparison of MF-score statistics of both desktop and (its counterpart) mobile versions of top sites.

	desktop (LP_d)	**mobile** (LP_m)
Min	0	0
Max	0.98	0.99
Median	0	0.58
Mean	0.209	0.487
Std. Dev.	0.306	0.322
IQR	0.43	0.605

1) Heuristics-based Ground Truth: First, we accessed the top-500 sites from mobile phone emulators, and identified URLs to their counterpart mobile version by following the re-direction of HTTP responses. At the time of repeated experiments, on average, 465 out of 500 top sites were accessible via HTTP requests, and 44.7% (208 out 465) of top sites turned out to maintain separate URLs/contents for mobile device users.From those 208 top sites that maintain two separate URLs/contents for desktop and mobile users, we consider LP_m (e.g., "google.com/m") as the candidate "ground truth" mobile-friendly landing pages. Our rationale is that since these LP_m landing pages are likely to be explicitly designed for mobile device users in mind by web designers of large top-500 companies with sufficient resources (as opposed to mom-and-pop kind of small businesses), their mobile-friendliness quality is deemed to be higher than that of LP_d.

To validate this, we perform the statistical hypothesis test, $H_0 : \mu_m - \mu_d = 0$, where μ_m and μ_d refer to the means of the populations LP_m and LP_d, respectively. Table I shows the basic MF-score statistics of those 208 top sites. Between two counterpart landing pages for the same company, for instance, mobile versions have higher average MF-score than desktop versions have (0.487 vs. 0.209), showing 133% improvement on average. With p-value < 0.01 (2-tailed t-test), finally, we reject H_0 and conclude that there exists a statistically significant difference between μ_m and μ_d, strongly implying that `mobileOK` is effective in differentiating the mobile-friendly vs. mobile-unfriendly landing pages.

2) Human-judged Ground Truth: Second, using only top-200 sites, three human judges visited each site using one of Apple iPhone 4, HTC Inspire 4G, and Samsung Galaxy S phones, and evaluated the site as either mobile-unfriendly (class=0) or mobile-friendly (class 1). Note that there was *no* previous agreement on the precise definition of mobile-friendliness among three judges. This way, we intended to simulate normal users' *ambiguous* perception on mobile-friendliness. Therefore, for instance, it is entirely possible that one judge regards a page with pop-up menus as class=0 while another as class=1. We then measured how much agreement among judges there is on the mobile-friendliness judgements

[4]http://aws.amazon.com/

[5]http://www.alexa.com/topsites/countries/US

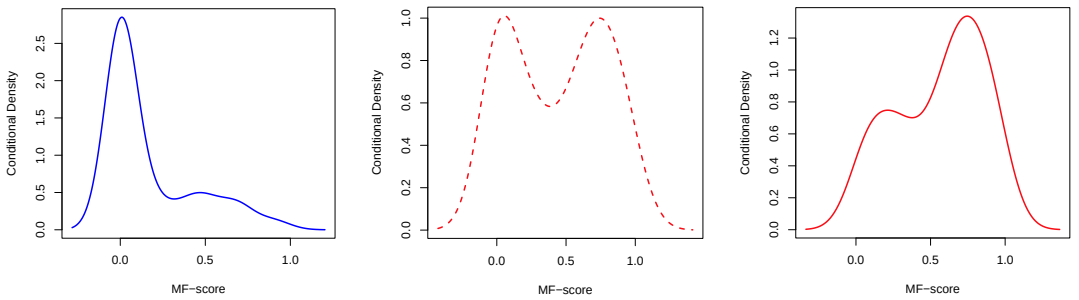

(a) mobile-unfriendly pages (class = 0)

(b) mobile-friendly pages (class = 1)

(c) mobile-friendly pages (class = 1 with MF-score > 0)

Fig. 1. Conditional probability distribution.

using the Cohen's *Kappa* measure from the social sciences, $Kappa = \frac{P(A)-P(E)}{1-P(E)}$, where $P(A)$ is the proportion of the observed agreement between two judges, and $P(E)$ is the proportion of the times two judges would agree by accident. Since there are three judges, we computed three pair-wise Kappa measures and used their average. In addition, since class (i.e., 0 or 1) distribution is skewed, we used the marginal statistics to calculate $P(E)$, and obtained a final Kappa value of 0.7126. In general, a Kappa value between 0.67 and 0.8 is regarded as fair agreement between judges. Therefore, we concluded that our human-judged ground truth is in a fair agreement, and decided to use the *majority voting* scheme–i.e., a site with at least two "1"s is labeled as class=1 and with at least two "0"s as class=0.

Next, after removing those 10 sites that were inaccessible, we obtained a total of 190 data points (i.e., 71 sites in class=0 and 119 sites in class=1). Finally, using `mobileOK`, we measured the MF-scores of all 190 sites. Figures 1(a) and (b) show kernel density estimators of MF-score distributions using data in class=0 and class=1, respectively. Observe that the high densities around 0 in Figure 1(a) and the bimodal shape around both 0 and 0.75 in Figure 1(b). This implies that when a page p gets a high MF-score, p's probability to be labeled as mobile-friendly by human judges is very high. However, when p's MF-score is close to 0, it could be either of two reasons: (1) MF-score does not well reflect the human perception on mobile-friendliness (since p that human judges viewed as mobile-friendly got the MF-score close to 0), or (2) p is simply poorly designed and mobile-unfriendly[6]. In other words, *the interpretation of cases with MF-score close to 0 should be made with care* (since it could mean one of two reasons). To further validate this implication, we removed 26 "contradicting" data points in class=1 whose MF-score is 0, and got Figure 1(c). Now, observe the unimodal distribution with high densities around 0.75. Figure 1(c) can be considered as an increasing function, implying that a page with a "higher" MF-score be "more" likely to be mobile-friendly.

[6]We noted that human judges cannot easily detect some potential drawbacks. For instance, page size is a critical factor in evaluating MF-score, yet people do not notice it unless they are in a poor mobile network.

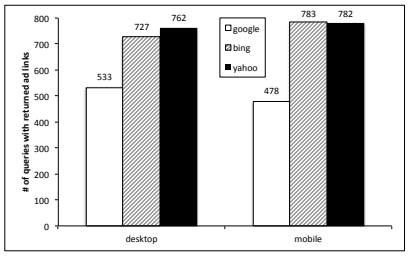

(a) Total # of queries (out of 800) that returned ad links

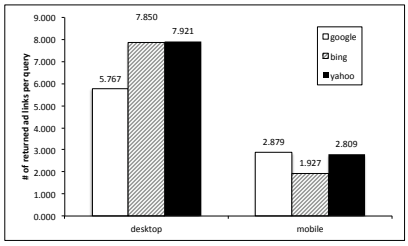

(b) Average # of returned ad links per query

Fig. 2. Basic statistics of ad links returned by 800 KDD Cup queries.

B. Basic Results

We next present the basic statistics of landing pages that we gathered. Figure 2(a) first shows the total number of queries (out of 800 KDD data set) that actually returned at least one ad link back. While both Bing and Yahoo share similar numbers (i.e., 727–783) for both desktop and mobile landing pages, Google has much fewer number of queries (i.e., 478–533) that returned at least one ad link. For instance, a query "`beacon federal`" to Google returns no ad links while that to Bing returns an ad link to a mortgage sale. Similarly, a query "`aircargo tracking`" to Yahoo or Bing returns ten ad links while zero to Google. Since we do not have an access to details of each search engine's ad placement algorithm, we cannot explain why such a difference exists. For all three search engines, between desktop and mobile landing pages, however, the numbers of queries that actually returned at least one ad link in sponsored search remain almost the same.

Figure 2(b) next shows the average number of ad links (in the first page) returned per query. Several things are noticeable here. First, on average, for desktop users, Google returns fewer

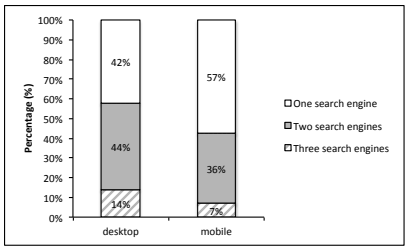

(a) Across 3 search engines

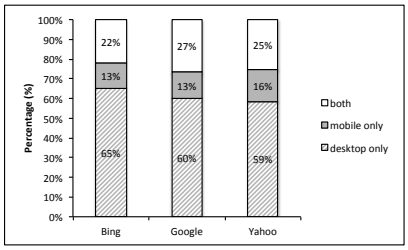

(b) Across 2 platforms (desktop and mobile)

Fig. 3. Distributions of ad links.

number of ad links (per query) than either Bing or Yahoo does (5.7 vs. 7.8 & 7.9). This is similar to the trend of Figure 2(a) where Google has fewer number of queries (out of 800) with at least one ad link returned. However, for mobile users, all three search engines returns similar number of ad links (per query)–i.e., 1.9–2.8. Second, overall, the average numbers of ad links for mobile landing pages are much smaller than those for desktop landing pages. Since real estates for mobile devices are more limited (e.g., smaller screen size), it is not practical to return the similar number of ad links as well as organic search results. Therefore, on the mobile platform, all three search engines appear to have reduced the number of ad links to return substantially.

Next, Figure 3 summarizes the distribution of ad links across 3 search engines and 2 platforms (i.e., desktop vs. mobile). We converted all URLs of ad landing pages to a simple host name format such as "*.foo.com" or "*.foo.org" and measure the overlap. For instance, both ad links "www.REI.com/Bikes" and "search.ib2.rei.com/?query=alpine" are converted to "*. rei.com". Note that this simple heuristic is not without errors. For instance, using the simple heuristics, we cannot identify that both URLs "www.microsoft.com/Xbox360" and "www. microsoftstore.com/Xbox_360" belong to the same company. Figure 3(a) shows that on the desktop, 42%, 44%, and 14% of companies bid for only one, two, and three search engines, respectively. Using the mobile data, on the other hand, 57% of companies bid for only single search engine. Since the total # of ad links returned for mobile devices is much smaller than that for desktop devices, the overlap across search engines on the mobile platform came lower. Overall, it shows that a substantial portion of companies participate in only one search engine's ad program, presenting ample opportunities for other search engines. Figure 3(b) shows that distribution of overlap of landing pages across both desktop and mobile platforms. All

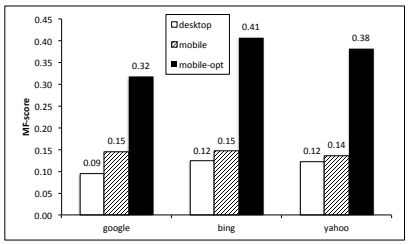

Fig. 4. Comparison of MF-score of 3 types.

three search engines show similar pattern. About 59%–65% of companies' landing pages are returned for only desktop devices while about 22%–27% of companies' landing pages are returned for both desktop and mobile devices.

C. Mobile-Friendliness of Landing Pages

Figure 4 summarizes MF-scores of three types of landing pages from three search engines: (1) `desktop` for LP_d, (2) `mobile` for LP_m, and (3) `mobile-opt` for the subset of LP_m whose URLs exhibit conventional patterns for mobile URLs such as "m.foo.com" or "foo.com/mobile/". Ideally, LP_m should include all landing pages that were specifically designed for mobile device users. However, in reality, advertisers often make one version of landing page (e.g., for desktop) and use them for both desktop and mobile users. Therefore, LP_m includes many landing pages poorly designed (although returned for mobile device users). Therefore, the `mobile-opt` type is our attempt to really filter out those mobile landing pages explicitly designed for mobile devices among LP_m. Table II shows the number of landing pages used in the measurement. The last column of Table II shows the percentage of the `mobile-opt` type within LP_m type.

Overall, the qualities (w.r.t. mobile-friendliness) of current landing pages, whether they were returned for desktop or mobile device users in sponsored search, are very poor–their MF-scores are only 0.09–0.15. In particular, we note that there is little difference in MF-scores between the `desktop` and `mobile` types. This implies that currently advertisers do not necessarily make their landing pages mobile friendly. This could be due to the lack of resources of such advertisers or unawareness of mobility standards. More interestingly, note that the MF-scores of `mobile-opt` type for all three search engines are significantly higher than MF-scores of the other two types are. That is, those mobile landing pages that are likely to be exclusively designed for mobile users indeed show much better mobile-friendliness than landing pages in the `desktop` and `mobile` types. However, in the `mobile` type, the percentage of `mobile-opt` type is still relatively small (i.e., see the last column of Table II). Therefore, we speculate that majority of landing pages returned for mobile users (i.e., the `mobile` type) are currently *not* mobile friendly yet (despite the existence of better mobile-friendly landing pages in the `mobile-opt` type). Current mobile landing pages are most likely created for "desktop" users in mind originally but being served to mobile users as well.

TABLE II
OF LANDING PAGES OF 3 TYPES.

Data set	desktop	mobile	mobile-opt (%)
google	2,860	1,258	203 (16.1%)
bing	4,915	1,388	192 (13.8%)
yahoo	5,072	2,013	260 (12.9%)

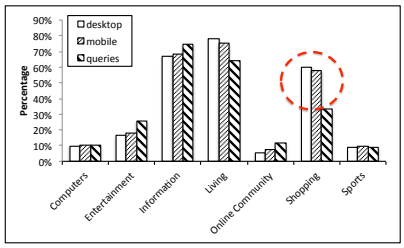

(a) Percentage of ad links per query type

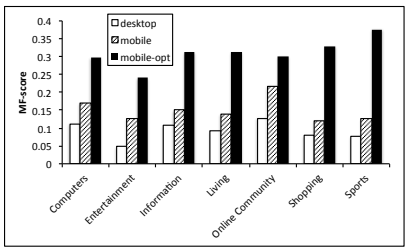

(b) MF-score of landing pages per query type

Fig. 5. Distribution among 7 query types returned from Google.

D. Classifying Landing Pages

Each query from 800 KDD Cup data set has a number of overlapping human-assigned query types assigned. The top-most level types include: *Computers*, *Entertainment*, *Information*, *Living*, *Online Community*, *Shopping*, and *Sports*. Since a query can have multiple types assigned, the sum of all percentages exceeds 100%. The distributions of query types and their returned ad link types are not balanced in general. In Figure 5(a), the third bar (labeled as "queries") represents % of query types. For instance, the *Information* and *Living* type queries are most frequent in Google data set (other graphs for Bing and Yahoo are similar and omitted). The first two bars (labeled as "desktop" and "mobile") show % of ad links per each query type on two platforms (e.g., an ad link returned for the *Shopping* type query is considered as the *Shopping* type ad link).

Note that, in the circled area of Figure 5(a), the percentages of ad links (i.e., first two bars) in *Shopping* category is noticeably higher than the percentage of query (i.e., third bar). That is, more ad links are returned for the *Shopping* type than others. This is probably so since shopping related queries should be more relevant to campaigned ads than other categories in sponsored search. Figure 5(b) shows the comparison of three types of landing pages among 7 query types. Overall, the same pattern as shown in Figure 4 (i.e., MF-scores of `mobile-opt` type landing pages are significantly higher than the other two types) is observed as well.

Becker et al [4] introduced 4 classes of landing pages:

TABLE III
BREAK-DOWN OF LANDING PAGES ACCORDING TO THE 4 CLASSIFICATION CLASSES.

Class	%HP	%ST	%CB	%O
Our %	32.62%	23.53%	29.41%	14.44%
% from [4]	25%	26%	37.5%	11.5%
AVG MF-score	0.2	0.12	0.098	0.281

(1) Homepage (`HP`) class landing page is the top-level page of an advertiser's web site; (2) Search Transfer (`ST`) class landing page is dynamically generated search results on an advertiser's web site; (3) Category Browse (`CB`) class landing page leads users to a sub-section of an advertiser's web site via navigation; and (4) Otherwise, the Other (`O`) class landing page. We first check the break-down of our landing pages according to these 4 classes. From the randomly chosen 200 landing pages in the Google `mobile` data set, 13 landing pages were removed since either they did not respond to HTTP requests or were not rendered at all. Then, we manually classified the remaining 187 landing pages into one of 4 classes, and the result is shown in Table III. Note that our break-down of 4 classes is not entirely in sync with that from [4]–ours tends to have more number of `HP` class landing pages and less number of `CB` class ones. With respect to average MF-score, `HP` class landing pages show roughly twice higher MF-score than `ST` or `CB` class landing pages do.

Next, in each class, we examined the impact of sub-tests inside `mobileOK` toward MF-score. From 25 tests, non-differentiating 6 tests (e.g., MAIN_DOCUMENT, CONTENT_FORMAT_SUPPORT, PROVIDE_DEFAULTS, STYLE_SHEETS_SUPPORT, and TABLES_ALTERNATIVES) are first removed (since all landing pages either passed or failed completely on those 6 tests). Figure 6 then shows the results of 4 classes using the remaining 19 tests. Y-axis denotes the *passing rate*–the fraction of landing pages that passed the particular test. Note that Figure 6(b) (for ST class landing pages) clearly has a different shape (as noted in three circles) compared to the rest of three graphs of Figures 6(a), (c), and (d). For instance, Figure 6(b) shows a rather low passing rate for AUTO_REFRESH while comparatively high passing rates for IMAGE_MAPS or TABLES_NESTED tests. Some possible explanations are below:

- AUTO_REFRESH: Auto-refreshing pages are widely recognized as one of the reasons to cause the web page accessibility problem. It is sometimes employed as a redirection mechanism. However, it also adds to delay on slow links. This is also the case in the ST type landing pages. Some ST type pages use <noscript> tag to redirect users to the final search result pages if the web browser does not support client-side scripting, or has disabled scripts in the browser setting. For instance, "ask.com" provides a slightly different search result page if it detects that the browser does not support scripting. Note that ST type landing pages are often dynamically generated to accommodate query keyword specific con-

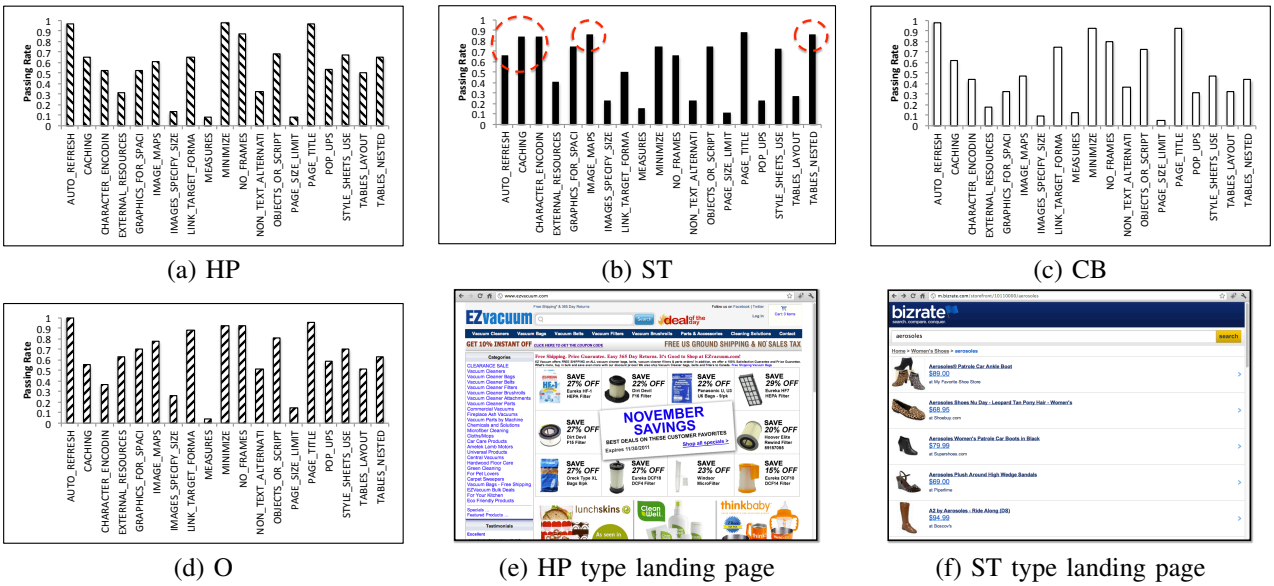

(a) HP (b) ST (c) CB

(d) O (e) HP type landing page (f) ST type landing page

Fig. 6. Passing rates break-down of 25 tests across 4 classes.

tents. In so doing, it appears that many ST landing pages are misusing auto-refreshing and thus get relatively low passing rates.

- TABLES_NESTED: Many web pages use tables for layout. When deeply nested tables are used in pages, however, it may cause rendering problems in mobile devices. Figure 6 shows that in all types of landing pages, the passing rates for TABLES_NESTED is not high (HP=0.66, ST=0.86, CB=0.44, and O=0.63). However, compared to others, in particular, ST type has a relatively higher passing rate of 0.86. Figures 6(e) and (f) show typical examples of HP and ST type landing pages. Often, ST type landing pages use a simple table layout (without much nesting) to display the contents of search results (e.g., Figures 6(f)), while other page types such as HP or CB often use complicated and nested table layouts to organize its content (e.g., Figures 6(e)).
- IMAGE_MAPS: In our pilot study, ST type landing pages show a higher passing rate for the test of IMAGE_MAPS. Other type landing pages often failed in this test since they used images for input element of buttons.

V. Findings and Implications

A. Current Landing Pages Are Not Mobile-Friendly Enough

The main finding of our study (esp. Figure 4) can be summarized as follows: "*Majority of current landing pages in sponsored search are not mobile-friendly*." Since the mobile-friendliness qualities of landing pages have a direct implication on the perception of users who click ad links, from search engine companies' point of views, it is critically important to keep their landing pages mobile-friendly. Similarly, advertisers also have keen interests in making their landing page mobile-friendly due to the explosive increase of access from mobile devices. Therefore, the fact that majority of mobile landing pages being "mobile-unfriendly" is somewhat surprising and problematic. To alleviate this issue, search engine companies that currently offer sponsored search and online advertising programs may attempt to help small-to-medium advertisers so that their landing pages become more mobile-friendly. This help could be an education program to raise the awareness of mobile-friendliness of landing pages for small advertisers.

TABLE IV
% OF LANDING PAGES THAT ARE GENERATED FROM TOP-500 MOST-VISITED US SITES OF ALEXA.COM.

Data set	`desktop`	`mobile`
google	20.7%	17.65%
bing	22.56%	23.49%
yahoo	19.79%	22.57%

Note that many current landing pages are from "small" businesses. In Table IV that shows the percentages of landing pages by top-500 most visited US sites of Alexa.com. the percentages range from 17.65% (for Google `mobile`) to 23.49% (for Bing `mobile`) and is roughly around 20% of all landing pages.

It is reasonable to assume that these top-500 companies are advertisers with sufficient resources (e.g., IT knowledge or personnel). Note that the percentages range from 17.65% (for Google `mobile`) to 23.49% (for Bing `mobile`) and is roughly around 20% of all landing pages. Unlike these 20% of large companies, therefore, the other 80% of *small* businesses may not have enough resources to make their landing pages optimized for mobile devices. Therefore, search engine companies (and related mobile communities) should develop more tools to aid the creation of mobile-friendly landing pages for small businesses.

B. Static Creation of Mobile Landing Pages Is Not Enough

Since keeping landing pages mobile-friendly is important for both advertisers as well as search engine companies,

commercial or open-source tools (e.g., WireNode, Mobify, bMobilized, Onbile, Google Page Speed) are in abundance to help advertisers create mobile-friendly web sites. Using such tools, businesses can either build mobile-friendly landing pages from the scratch or convert existing desktop-oriented landing pages to mobile-oriented ones. By and large, these software solutions use some forms of heuristics (possibly based on common practices of building mobile web pages or W3C's standards) and generate results *statically*–i.e., for an input web page, a solution generates always the same output web page. Recall that the definition of "mobile-friendliness" is inherently subjective and dynamic. Software/hardware characteristics of mobile devices (e.g., screen size, CPU, barebone vs. smartphone) and network infrastructure (e.g., Mobile WiMax, 3GPP Long Term Evolution, 4G) are diverse and rapidly changing. Similarly, types and features of landing pages are galore. Therefore, it is hard to build a generic solution that works for *all* variations. Because of this reason, existing software solutions determine mobile-friendliness statically based on common practices. However, ultimately, the reason for advertisers to have mobile-friendly landing pages is because: "*advertisers want users to experience the same interactions across different devices*."

When an advertiser has two landing pages, p for desktop and q for mobile users, if she witnesses the equal traffic patterns and conversion rates between p and q, then her ultimate goal is satisfied (and q can be viewed as a mobile-friendly version of landing page of p, regardless of its MF-score). Taking this insight, suppose there is a mechanism, M, to be able to closely track user experiences embedded into each landing page. Then, consider a landing page, p, for desktop users and $M(p)$, a set of statistical features derived from the mechanism M embedded in p. Then, the problem of making mobile-friendly landing pages can be re-casted as follows:

Problem 1 *Given a desktop landing page p with $M(p)$, generate a mobile landing page q such that the probability $P(M(p) = M(q))$ is high.* ∎

The probability $P(M(p) = M(q))$ becomes 1 when user interactions against p and q are identical–i.e., best scenario for advertisers, and becomes 0 when the bounce rate of q becomes 1 and all users landing on q immediately exit the site–i.e., worst scenario for advertisers. Naturally, this probability is a good objective measure to indicate how good the rewriting of p into q is. We believe that this notion of data-driven mobile-friendliness and its mapping to MF-score is entirely novel and desired. We leave the interesting research issue of how to enable this data-driven mobile-friendliness to tools like `mobileOK` as future work.

VI. Conclusion

We have analyzed a total of 60,419 real ad links collected from three search engines with the focus on the mobile-friendliness of landing pages. From our study, we have found: (1) The mobile-friendliness of a given web page can be effectively measured by W3C's `mobileOK`. As such, landing pages with a higher MF-score are likely to be rendered better in mobile devices; (2) Substantial portions of advertisers currently participate in only one search engine's ads program (i.e., 42–57%), presenting further marketing opportunities for other search engines or online ads programs; and (3) Majorities of currently-serving landing pages on both desktop and mobile platforms have very poor MF-scores, and are *not* likely to be rendered well in mobile devices. Since the usage of mobile web browsing increases rapidly, to attract more incoming traffic (and thus increased conversions), businesses should put more efforts into building mobile-friendly web sites and landing pages.

Based on our findings, we also proposed a novel notion of mobile-friendliness that is based on collected user experience data, instead of a fixed set of heuristics. We plan to investigate how to incorporate such a notion of *data-driven mobile-friendliness* in measuring MF-score of landing pages, converting desktop-oriented landing pages to mobile-oriented ones, or verifying the equivalence between desktop-oriented and mobile-oriented landing pages (w.r.t. contents, functionalities, and user experience).

Acknowledgment

This research was in part supported by NSF DUE-0817376 and DUE-0937891 awards, and Amazon web service research grant (2011).

References

[1] The Mobile Internet Report. Tech. rep., Morgan Stanley Research, Dec. 2009.

[2] Arase, Y., Hara, T., and Nishio, S. Web Page Adaptation and Presentation for Mobile Phones. In *Handheld Computing for Mobile Commerce: Applications, Concepts and Technologies*. 2010, 240–262.

[3] Baumgartner, R., Fayzrakhmanov, R., Gattringer, R., Gbel, M., Holzinger, W., Klein, D., and Kruepl, B. Web 2.0 Vision for the Blind. In *ACM Web Science Conf.* (Apr. 2010).

[4] Becker, H., Broder, A., Gabrilovich, E., Josifovski, V., and Pang, B. What Happens after an Ad Click? Quantifying the Impact of Landing Pages in Web Advertising. In *ACM CIKM* (Nov. 2009).

[5] Chen, Y., Xie, X., Ma, W.-Y., and Zhang, H.-J. Adapting Web Pages for Small-Screen Devices. *IEEE Internet Comp. 9* (Jan. 2005), 50–56.

[6] Choi, Y., Fontoura, M., Gabrilovich, E., Josifovski, V., Mediano, M., and Pang, B. Using Landing Pages for Sponsored Search ad Selection. In *Int'l Conf. on World Wide Web (WWW)* (2010).

[7] Erol, B., Berkner, K., and Joshi, S. Multimedia Clip Generation From Documents for Browsing on Mobile Devices. *IEEE Trans. on MultiMedia (TMM) 10* (Aug. 2008), 711–723.

[8] Hu, W., Kaabouch, N., Yang, H., and Hu, W. Technologies and Systems for Web Content Adaptation. In *Handheld Computing for Mobile Commerce: Applications, Concepts and Technologies*. 2010, 263–277.

[9] Li, Y., Zheng, Z., and Dai, H. K. KDD CUP-2005 Report: Facing a Great Challenge. *SIGKDD Explor. Newsl. 7* (Dec. 2005), 91–99.

[10] Purwandari, B., Hall, W., and De Roure, D. The Impact of the Mobile Web in Developing Countries. In *ACM Web Science Conf.* (Apr. 2010).

[11] Purwandari, B., Hall, W., and Wills, G. Methodology for Impact Analysis of the Mobile Web in Developing Countries: a Pilot Study in Nairobi, Kenya. In *ACM Web Science Conf.* (Jun. 2011).

[12] Ravi, S., Broder, A., Gabrilovich, E., Josifovski, V., Pandey, S., and Pang, B. Automatic Generation of Bid Phrases for Online Advertising. In *ACM Int'l Conf. on Web Search and Data Mining (WSDM)* (2010).

[13] Xiao, X., Luo, Q., Hong, D., Fu, H., Xie, X., and Ma, W.-Y. Browsing on Small Displays by Transforming Web Pages into Hierarchically Structured Subpages. *ACM Trans. on the Web (TWEB) 3* (Jan. 2009).

[14] Zhang, D. Web Content Adaptation for Mobile Handheld Devices. *ACM Comm. of the ACM (CACM) 50* (Feb. 2007).