What Happens after an Ad Click? Quantifying the Impact of Landing Pages in Web Advertising

Hila Becker[‡1], Andrei Broder[†], Evgeniy Gabrilovich[†], Vanja Josifovski[†], Bo Pang[†]

‡ 1214 Amsterdam Avenue, New York, NY 10027, USA
† Yahoo! Research, 2821 Mission College Blvd, Santa Clara, CA 95054, USA
hila@cs.columbia.edu | {broder | gabr | vanjaj | bopang}@yahoo-inc.com

ABSTRACT

Unbeknownst to most users, when a query is submitted to a search engine two distinct searches are performed: the *organic* or *algorithmic* search that returns relevant Web pages and related data (maps, images, etc.), and the *sponsored search* that returns paid advertisements. While an enormous amount of work has been invested in understanding the user interaction with organic search, surprisingly little research has been dedicated to what happens *after* an ad is clicked, a situation we aim to correct.

To this end, we define and study the process of *context transfer*, that is, the user's transition from Web search to the context of the *landing page* that follows an ad-click. We conclude that in the vast majority of cases the user is shown one of three types of pages, namely, *Homepage* (the homepage of the advertiser), *Category browse* (a browse-able subcatalog related to the original query), and *Search transfer* (the search results of the same query re-executed on the target site). We show that these three types of landing pages can be accurately distinguished using automatic text classification. Finally, using such an automatic classifier, we correlate the landing page type with *conversion* data provided by advertisers, and show that the conversion rate (i.e., users' response rate to ads) varies considerably according to the type. We believe our findings will further the understanding of users' response to search advertising in general, and landing pages in particular, and thus help advertisers improve their Web sites and help search engines select the most suitable ads.

Categories and Subject Descriptors

H.3.m [**Information Storage and Retrieval**]: Miscellaneous

[1]The research described herein was conducted while the first author was a summer intern at Yahoo! Research.

CIKM'09, November 2–6, 2009, Hong Kong, China.

General Terms

Experimentation, Measurement

Keywords

online advertising, landing page taxonomy

1. INTRODUCTION

In recent years, online advertising has become a prominent economic force that sustains numerous Internet services, ranging from major Web search engines to obscure blogs. The standard approach to textual Web advertising is based on modeling the user's needs and interests to find suitable ads. In particular, in Web search, numerous studies have focused on classifying the query intent [2, 5, 9, 16] and on retrieving the most relevant ads [3, 12, 14]. However, surprisingly little research has been devoted to what actually happens *after* an ad is clicked.

The ultimate goal of advertising is *conversion*, that is, the transformation of a consumer who has noticed the ad into a buyer of the product or service being advertised. Here, "buyer" should be construed in a general sense: in a political campaign, a "buy" is a vote for the candidate; for a car advertiser, a "buy" might be a test-drive at the dealership; for an on-line publication or service, a "buy" might be a free subscription, etc.

In this paper we focus on *sponsored search* advertising, which displays textual ads alongside algorithmic (or *organic*) search results. In this case the Web search query issued by the user embodies the quintessence of their intent, and is the main trigger for selecting ads to display. Once the search engine result page is presented, a user potentially becomes a "buyer" in two stages:

1. **Clickthrough** First the user must click on the ad displayed in response to their query. As a result, the user is transferred to the *landing page* for this (query, ad) combination, which is defined as the first page the user sees on the advertised Web site. Usually advertisers pay the search engine for every click on their ads — this is the *cost-per-click* or *CPC* model (see [7] for more details). The observed frequency with which a particular ad is clicked for a particular query is called the *clickthrough* rate, CTR(query, ad).

2. **Conversion** At this stage, the user, possibly after a certain amount of activity on the advertiser's site, be-

comes a "buyer" of the product or service being advertised. The observed frequency with which clickers on a given ad become "buyers" is called the *conversion* rate. In some cases, the advertisers pay only for conversions. To emphasize that "conversion" can be a generic action, not just a conventional buy, this is called the *cost-per-action* or *CPA* model.

Understanding the conversion rate is essential for search engines and advertisers. In the CPC model, it determines the advertisers' return on investment and informs the search engines about the value of their product; in the CPA model, it determines directly how much money changes hands. To this end, we define and study the process of *context transfer*, that is, the user's transition from her previous activity (to wit, Web search) to the different possible contexts found on the landing page after clicking on an ad. Arguably, a careful choice of the type of context transfer is among the most important factors explaining the subsequent conversion. We introduced the basic concept of context transfer in a previously published poster [1]. In this paper we expand our experiments and further explore the significance of context transfer by studying the correlation between the type of context transfer and the observed conversion rate.

After reviewing a comprehensive sample of several hundred ads and corresponding landing pages, we found that the vast majority of the observed context transfers fell into one of the following three classes:

1. **Homepage** Here the landing page is simply the home page of the advertiser's Web site. This can be appropriate both for small mom-and-pop businesses, which cannot afford or do not need more sophisticated structures, and for large online stores, which usually populate their homepage with daily promotions in addition to describing the variety of their offerings.

2. **Category browse** Here the landing page is a browseable sub-catalog of products being offered on the advertiser's site. This is usually suitable for queries related to a meaningful group of products. For example, an ad shown for the query "California Zinfandel" can have a landing page devoted to a variety of Zinfandel wines (see Figure 1(b)).

3. **Search transfer** In this case, clicking on an ad leads to the results of a search conducted on the advertiser's Web site using the original query that triggered the ad (or slightly modified versions of it). This context transfer is suitable when a query has multiple interpretations or is relevant to numerous offerings, or the target Web site does not have a corresponding category (see Figure 1(a)).

We observed that these three classes combined account for over 88% of the ads in our sample dataset. Furthermore, these classes are easily distinguishable and we were able to build a high accuracy ($> 80\%$) classifier for them. Using this classifier, we then conducted a study of correlation between the different types of landing pages and the conversion rates of the corresponding ads, when available to us. (Advertisers sometimes provide conversion data to search engines; see further discussion of the conversion dataset in Sections 3 and 5.) Our final results are based on over 30,000 unique landing pages, automatically classified.

We also examined the suitability of different classes of landing pages for different types of queries (e.g., queries of different lengths or on different topics). Interestingly, in our dataset there seems to be little agreement among advertisers as to which landing page to use for which query, as for many query types we observed actual use of a wide variety of landing pages. However, we found that in many cases the existing choice of landing pages could be sub-optimal, and we encourage advertisers to experiment with different types of landing pages, and then make an informed choice based on statistical evidence.

The contributions of this paper are threefold. First, we propose a taxonomy of ad landing pages. Second, we use standard machine learning techniques to build a classifier capable of automatically mapping landing pages onto the classes of this taxonomy. Finally, we juxtapose the frequency of actual use of different classes with their reported conversion rates. Based on our findings, we encourage advertisers to conduct principled studies of the effect of different classes of landing pages on conversion rates.

From a scientific perspective, the idea that advertising is a form of information has been promulgated for over 30 years [11]. However, the challenge of *retrieving* this type of information has become pertinent only with the advent of Web advertising and the practical necessity of choosing the "best" among millions of competing ads. Many of the proposed solutions are based on classic IR methods and were the subject of several information retrieval papers in recent years (see e.g. [3, 12, 14]). In this context, our study aims to illuminate one aspect in which Web advertising information differs from both classical information (documents) and non-interactive advertising information—namely, in Web advertising, the information creator has significant control on how this information is used by its consumer.

2. BACKGROUND AND RELATED WORK

We begin by providing some background on the field of computational advertising, and then discuss relevant related work.

Background: Textual advertising on the Web.

A large part of the Web advertising market consists of *textual ads*, the ubiquitous short text messages usually marked as "sponsored links." There are two main channels for distributing such ads. *Sponsored search* (or *paid search advertising*) places ads on the result pages of a Web search engine, where ads are selected to be relevant to the search query (see [7] for a brief history of the subject). All major Web search engines support sponsored ads and act simultaneously as a Web search engine and an ad search engine. *Content match* (or *contextual advertising*) places ads on third-party Web pages. In this paper we focus on sponsored search. However, we believe that the taxonomy of landing pages we propose here could be easily adapted for modeling conversion rates in the content match scenario, and plan to investigate this direction in future work.

Sponsored search is an interplay of three entities. The **advertiser** provides the supply of ads. Usually the activity of the advertisers is organized around *campaigns*, which are defined by a set of ads with a particular temporal and thematic goal (e.g., sale of digital cameras during the holiday season). As in traditional advertising, the goal of the advertisers can be broadly defined as promotion of products or services. The **search engine** provides "real estate" for placing ads (i.e., allocates space on search results pages), and selects ads that are relevant to the user's query. **Users** visit the Web pages and interact with the ads.

Sponsored search usually falls into the category of *direct marketing* (as opposed to *brand advertising*), that is, advertising whose aim is a "direct response," where the effect of a campaign is measured by the user reaction (e.g., purchase of advertised goods or services). Compared to the traditional media, one of the advantages of online advertising in general and sponsored search in particular is that it is relatively easy to measure the user response. Usually the desired immediate reaction is for the user to follow the link in the ad and visit the advertiser's Web site. However, the desired eventual outcome is for the user to perform a transaction on the advertised Web site, e.g., purchase a product or service being advertised. Therefore, our evaluation methodology is based on measuring *conversion rate*, which is the fraction of users who performed the desired transaction among those who merely clicked on the ad

The prevalent pricing model for textual ads is that the advertisers pay for every click on the advertisement (pay-per-click or PPC). The amount paid by the advertiser for each click is usually determined by an auction process [6], where the advertisers place *bids* on a search phrase. Thus, each ad is annotated with one or more *bid phrases*. In addition, an ad also contains a *title*, a *creative* (a few lines of textual descriptions), and a URL to the advertised Web page, called the *landing page*.

In the model currently used by all the major search engines, bid phrases serve a dual purpose: they explicitly specify queries that the ad should be displayed for and simultaneously put a price tag on a click event. Obviously, these price tags could be different for different queries. For example, a contractor advertising his services on the Internet might be willing to pay a small amount of money when his ads are clicked from general queries such as "home remodeling," but higher amounts if the ads are clicked from more focused queries such as "hardwood floors" or "laminate flooring." Most often, ads are shown for queries that are expressly listed among the bid phrases for the ad, thus resulting in an *exact match* (i.e., identity) between the query and the bid phrase. However, it might be difficult (or even impossible) for the advertiser to list all the relevant queries ahead of time. Therefore, search engines can also analyze queries and modify them slightly in an attempt to match pre-defined bid phrases. This approach, called *broad* (or *advanced*) match, facilitates more flexible ad matching, but is also more error-prone, and only some advertisers opt for it.

There are two bodies of prior research that are relevant to our study:

Online advertising. Online advertising is an emerging area of research, so the published literature is quite sparse. A recent study [17] confirms the intuition that ads need to be relevant to the user's interest to avoid degrading the user's experience and increase the probability of reaction.

There are several models of pricing online ads, which vary by the amount of risk shared by the advertiser and the publisher. Charging advertisers for ad displays (impressions) effectively places all of the risk with the advertiser, since the ads displayed might not even be relevant to the user. Charging in proportion to the conversion rate, which measures the proportion of users who actually committed to the advertised transaction, moves the risk almost entirely to the advertiser. Although many users perform a purchase in the same session when they click on the ad, many others will do so at a later time, having considered the worthiness of the transaction and conducting some research. In such cases, it becomes nearly impossible to relate the transaction to the initial ad click, making it very difficult to charge commensurately to the true conversion rate. The current practice of charging per click offers a middle ground between these two extremes, as paying per click lets the advertiser ascertain that the ad was at least somewhat relevant to the user, who expressed some interest by clicking on the ad

Due to this prevalence of charging per click, prior studies on forecasting users' response to ads mostly focused on predicting the click-through rates based on estimated ad relevance as well as click history [13, 15]. In contrast, in this work we study the true conversion rate.

Understanding user goals. Another relevant area of prior research focused on characterizing users' goals and information needs. Broder [2] formulated a taxonomy of Web search queries, which correspond to different types of users' information needs. Several subsequent papers also studied users' goals in Web search, notably [5, 8, 9, 16]. However, they mostly focused on characterizing the process of finding Web sites that satisfy the user's information need. In this work, we propose a taxonomy of landing pages for online advertising, which characterizes different scenarios of users' interaction with advertised Web sites.

3. DATASETS

We strove to define the taxonomy that is both concise and general enough to cover the majority of landing pages observed in real-life datasets. In this section, we describe the datasets used in this study and motivate their choice, and in the next section we proceed with the development of the taxonomy.

Clearly, the choice of sampling techniques used to form the datasets is crucial, as it affects the interpretation of the results. We created three datasets representing different underlining distributions of ads. All the datasets described were obtained from Yahoo! Web Search.

Pilot dataset: A small set of 200 unique sponsored search landing pages, which we used to define the taxonomy of landing pages and to construct an automatic landing page classifier. These landing pages belong to advertisements that were triggered by issuing 200 unique queries to Yahoo! Search. The queries were randomly sampled out of the 800 labeled queries used for the 2005 KDD Cup [10]. We used stratified sampling, dividing the set of KDD Cup queries into deciles according to query frequency computed from Web search query logs, and sampling 20 queries uniformly from each decile. Thus, this dataset was constructed to represent ads that are *shown* for both popular and rare queries.

Conversion dataset: Over 31,000 unique pairs of queries and landing page URLs, attributed with conversion information for one month in 2008, provided by participating advertisers. The conversion data was collected by adding http redirects to the links on the advertiser's site that represent conversion events (e.g., a "Buy" button). We used this dataset to validate our taxonomy definition, as well as to analyze the correlation between different types of landing pages and the corresponding conversion rates.

Browsing dataset: Actual conversion data is not always available, as many advertisers choose not to report it to the search engine. We define a proxy for conversion rate by using activity logs collected from a browser toolbar plug-in, which correspond to search trails starting with users' clicks on sponsored search results (see Section 5.3 for more details). The browsing dataset consists of over 66,000 landing pages as well as subsequent visits to other pages on the same site. This dataset represents a less biased sampling of clicked ads as it is not restricted by participation from advertisers.

4. TAXONOMY OF LANDING PAGES

In this section, we discuss the taxonomy of sponsored search landing pages. We start by describing a study that we conducted, which led to the initial definition of the landing page taxonomy. We then discuss the different classes of landing pages, outlining how advertisers transfer context by selecting to display a particular type of landing page. Finally, we describe a classifier built to automatically label a large set of landing pages, which we used in subsequent analysis.

4.1 Pilot study: defining landing page types

We began by conducting a study to test the feasibility of defining a landing page taxonomy. We wanted to observe whether or not sponsored search landing pages fall into a natural, unambiguous set of classes that could be easily characterized and identified by a human judge. For this purpose, we used the pilot dataset described in Section 3. We inspected each landing page in isolation, noting its structure, appearance and functionality. We observed several distinct, non-overlapping classes that sponsored search landing pages fall into. Each class represents a different context transfer technique that transitions the user from the search engine result page to the advertiser's landing page. It is interesting to note how much or how little context the advertiser preserves by using each class of landing pages.

Homepage *(HP).* This is the top-level page of the advertiser's Website. Many advertisers choose to simply display their home page as a landing page for their ads, often regardless of the query that triggered the ad. As noted, this approach is commonly used by either smaller, less experienced advertisers or well known brand-name advertisers that display their homepage when bidding on brand keywords. This approach might also be convenient for advertisers that bid on hundreds or thousands of keywords and do not want to make the investment to create a specific page for each keyword. Unless the user searched for the advertiser's brand name, using the homepage as a landing page does not make for a strong context transfer. For instance, consider a search for the word "Toyota." If Toyota is the advertiser, directing the searcher to Toyota's homepage will likely satisfy the user's information need. On the other hand, any other advertiser that does not have a Website dedicated to Toyota cars, e.g., a Website that provides price quotes for all car makes, would lose some of the context by showing a generic homepage, which does not immediately satisfy the search query (even though the relevant content may be found on the advertiser's Web site by following hyperlinks).

Search transfer *(ST).* Landing pages of this type are dynamically generated search results on the advertiser's site. This is a situation where the advertiser uses the original Web search query (sometimes with slight modifications) to perform a new search within its own site, and displays the results as the ad's landing page. For example, given a query "California Zinfandel," an online wine store would return a landing page similar to Figure 1(a), dynamically displaying search results. In landing pages of this type, context transfer is very strong only if the query used to generate the search results corresponds to products, services or information that the Website actually offers. However, many advertisers that use this technique do not design their campaigns carefully enough to ensure that all phrases they bid on yield meaningful search results, in which case the context is completely lost. On the other hand, this approach, similar to Homepage, does not require the investment to create a specific page for each keyword or group of keywords.

Category browse *(CB).* A Category browse landing page leads the user to a sub-section of the Website that is generally related to the query. This page is not at the top level of the Website (homepage) but rather could be navigated to from other pages on the site. Let us continue the previous example of an online wine store advertising for the query "California Zinfandel." Here, a Category browse landing page might describe the Zinfandel section of the Web site (Figure 1(b)). A small number of pages in our dataset described a single specific product. For convenience, we include them in the Category browse class, since from an user point of view, they also represent a transition from a searching activity to a query-specific browsing activity.
Category browse is a technique that advertisers can use both if the bid phrase refers to a general class of products or services, or to a specific one. If the user is looking for a general class of products, choosing a Category browse landing page would bring them one step closer to the product they are searching for. If the user is looking for a specific product, while the advertiser only carries different but related products, showing a category page allows the advertiser to present such related offerings. Compared to Homepage and to Search transfer, the Category browse technique requires more investment to create or identify a specific page for each keyword or group of keywords.

Other *(O).* These are standalone pages that appear to be disconnected from the rest of the Web site. These pages generally do not have many outgoing links and there is no way to reach them from the home page. One example of this class are standalone forms, where the sole purpose of the page is to gather information from the user. Another example is promotion pages, which supply promotional information about a product or service. These pages are similar to print ads in a newspaper, and often include phrases such as "try it now," "limited time," and "special offer."

4.2 Distribution of landing page types

We labeled each landing page in the pilot dataset according to the classes described in the previous section. The first line of Table 1 presents the distribution of labels. Note that the first three classes combined account for over 88% of the ads in this sample.

(a) Search transfer

(b) Category browse

Figure 1: Example landing pages

Class	%HP	%ST	%CB	%O
All	25	26	37.5	11.5
Info. companies/industries	24.6	28.4	37.7	9.3
Shop. stores/products	19.2	28.7	45.5	6.6
Shop. guides/research	20.9	27.1	47.3	4.7
Info. local/regional	29	16.9	33.9	20.2

Table 1: Distribution of page types in the pilot dataset with breakdown for sample query classes

Since the queries in our study were sampled out of the manually classified set for the KDD Cup, we were able to analyze our data with respect to the provided classes. We used an aggregate of the labels assigned by three human judges. See Table 1 for a breakdown of landing page types for the four most frequent query classes.

The distributions for each query class roughly follows the overall distribution, with no clear query-intent-based optimization. In particular, it is interesting to note that the breakdown of landing page types for queries in categories “Shopping: Buying Guides & Researching” and “Shopping: Stores & Products” follows a similar trend. Intuitively, if an advertiser knows that the user is researching a product, an appropriate strategy might be to use the home page in order to promote brand awareness. On the other hand, when the shopping intent is clearly focused on specific products and stores, one would assume that a more focused Category browse or even Search transfer page would be more appropriate.

4.3 Landing page classifier

In order to make meaningful claims about the impact of our findings, we need to obtain a larger set of landing pages and label them according to the taxonomy. We were able to train a sufficiently accurate classifier using standard machine learning techniques, which we briefly describe below.

To train the landing page classifier we used the set of landing pages with four-way labels annotated in the pilot study (Section 4.1). The features we used include the traditional bag-of-words representation of the visible landing page text with simple *tf.idf* weights, and the number of occurrences of frequently observed HTML patterns that appeared in the

Class	Precision	Recall	F-Measure
Homepage	0.917	0.786	0.846
Search transfer	0.862	0.926	0.893
Category browse	0.645	0.87	0.741
Other	0.5	0.25	0.333

Table 2: Performances of landing page type classifier on the test data.

landing page’s HTML source (e.g., a list of links separated by the characters ’>’ or ’:’, which usually appear on Category browse pages). Other useful features to note include the percentage of HTML overlap between the landing page and its base URL (which can help identify Homepage landing pages), and the ratio of form elements to text (as high ratio is commonly found in Other page types). We trained a Support Vector Machine model using Weka’s SMO implementation [18] on the reduced feature space induced from a supervised attribute selection technique, aiming to optimize the accuracy of the most frequent classes that accounted for more than 88% of the data. With 10-fold cross validation on the training data, our classifier accurately predicted the class label for 83% of the examples.

To ensure that our model is not overfitting, we tested our classifier on a separate test set that consisted of 100 manually labeled landing pages sampled from the browsing dataset, Accuracy of the classifier on the test set is 80%; Table 2 presents a breakdown of the performance by class.

5. CONVERSION OF LANDING PAGES

Conversion is at the core of the value-add generated by the search engine for all sponsored search participants. It is the ultimate goal of the advertisers: their return on investment in sponsored search depends directly on the conversion brought by the ads placed in the sponsored search systems. For a user, a conversion is an indication that the user has satisfied the intent of the query. Satisfied advertisers and users would make the search engine’s business model more viable by increased bids and more opportunity to earn revenue.

The taxonomy of landing pages proposed in Section 4 and the automatic detection of these landing page types can facilitate analysis of user behaviors after a click on a sponsored search ad. In this section, we examine the relationship between the landing page types and conversion rates.

5.1 Conversion rate

We define a *conversion* as a visit where the user performs the desired action, which can take many different forms ranging from further browsing, user registration, to product sales. For a given landing page URL (u) in an ad campaign, *conversion rate* ($cr(u)$) is the percentage of visitors who took the desired action, i.e., the ratio between the number of conversions and number of clicks associated with u.

For this study, we report the average conversion rate for a group of URLs ($u \in U$). One possibility is to define average conversion rate using the unweighted average conversion rate of all URLs, treating each URL equally, regardless of the number of clicks it received ($click(u)$). Since the conversion rates of URLs with more clicks provide more reliable estimates than the conversion rates of URLs with fewer clicks, we define the average conversion rate of U as the weighted average over $cr(u)$.

$$\text{avg. cr (U)} = \frac{\sum_{u \in U} cr(u) * log(click(u))}{\sum_{u \in U} log(click(u))},$$

and report the relative average conversion rate:

$$\text{rel. avg. cr (U)} = \frac{\text{avg. cr (U)} - \text{avg. cr (D)}}{\text{avg. cr (D)}},$$

where D denotes the entire dataset.

Note that we use the log function to scale down $click(u)$ before taking the weighted average – using $click(u)$ directly would have the undesirable effect of letting the conversion rates of popular landing pages completely dominate the average. This measure also effectively ignores the conversion rates of URLs with only one click. While it is possible to define a modified weight function to avoid this, we considered it reasonable to exclude URLs with too few clicks and used this measure as-is.

5.2 Correlation study with conversion dataset

In this section, we examine whether there is any correlation between the type of landing page used and the corresponding conversion rate. First, we examine the dataset in more detail.

Conversion dataset. As noted in Section 3, we obtained conversion information from participating advertisers. To facilitate our analysis, this dataset was augmented with additional information, removing entries with missing information in the process. For each landing page URL u and the query q that led to a visit to u, we collected:

- Number of clicks on u
- Number of conversions associated with u
- Price: average price paid to the search engine for each click on the query that led to u.
- Landing page type: we crawled the landing page and applied our automatic landing page type classifier when the text content of the page was non-empty.
- Query frequency: frequency[1] in Web search log for q.
- Query class: Optionally, we also included the class label of the query predicted by an automatic query classifier with respect to a commercial taxonomy of over 6000 nodes [4].

[1] Note that only queries appearing at least six times were retained in the dataset.

This resulted in a dataset of over 31,000 unique pairs of queries and landing page URLs. A tally of the query class labels predicted for each q revealed that our dataset covered a broad range of topics.

5.2.1 The overall picture

Table 3 summarizes the overall breakdown of different types of landing pages in the conversion dataset, and the relative average conversion rate associated with each type. As we can see, Category browse and Search transfer classes are

Class C	HP	ST	CB	O
Distribution	13.7%	33.7%	44.8%	7.8%
rel. avg. cr(C)	1.00	-0.55	-0.15	1.04

Table 3: Classifier class distribution and relative average conversion rate in the conversion dataset

the dominant choices, although the average conversion rates for them are lower than the average of the entire dataset. This does not necessarily imply that advertisers are choosing the wrong types of landing pages. Rather, these results point out that on average, advertisers who choose the Other and Homepage landing page types tend to have higher relative conversion rates than those who choose Category browse and Search transfer landing pages. One reason for this may be tied to the advertisers' varying definitions of conversion. The Other landing page type often contains stand-alone forms where a conversion may be the form's submission. On the other hand, a Search transfer landing page usually displays a list of products, where a conversion may correspond to a product sale. Clearly it is more difficult to achieve conversion in the latter case. Hence, we do not claim here that the choice of landing page type is the only factor that affects conversion. Instead, we provide analysis and insight into the correlation of landing pages and conversion.

With these caveats in mind, we proceed to explore the correlation between landing page types and conversion rates for different types (groupings) of queries.

5.2.2 Analysis of different groups of queries

We start by examining the landing page type usage and conversion information for different query frequencies (Figure 2(a) & 2(b)), different query lengths (Figure 2(c) & 2(d)), as well as different prices paid (Figure 3(a) & 3(b)) and different query classes (Figure 3(c) & 3(d)). One consistent trend is that the Other class is the least frequently used landing page type, with the highest or the second highest average conversion rate. As we discussed earlier, since the Other class includes registration pages and the like, the conversions may be less comparable. Thus, we focus our analysis on the three dominant classes.

Overall we observe similar trends as seen on the entire dataset: Category browse and Search transfer classes are used more often, but typically achieve lower conversion rates. Furthermore, the relative orderings in terms of both usage and conversion are mostly consistent, regardless of the topics (i.e., query class labels) of the queries (Figure 3(c) and

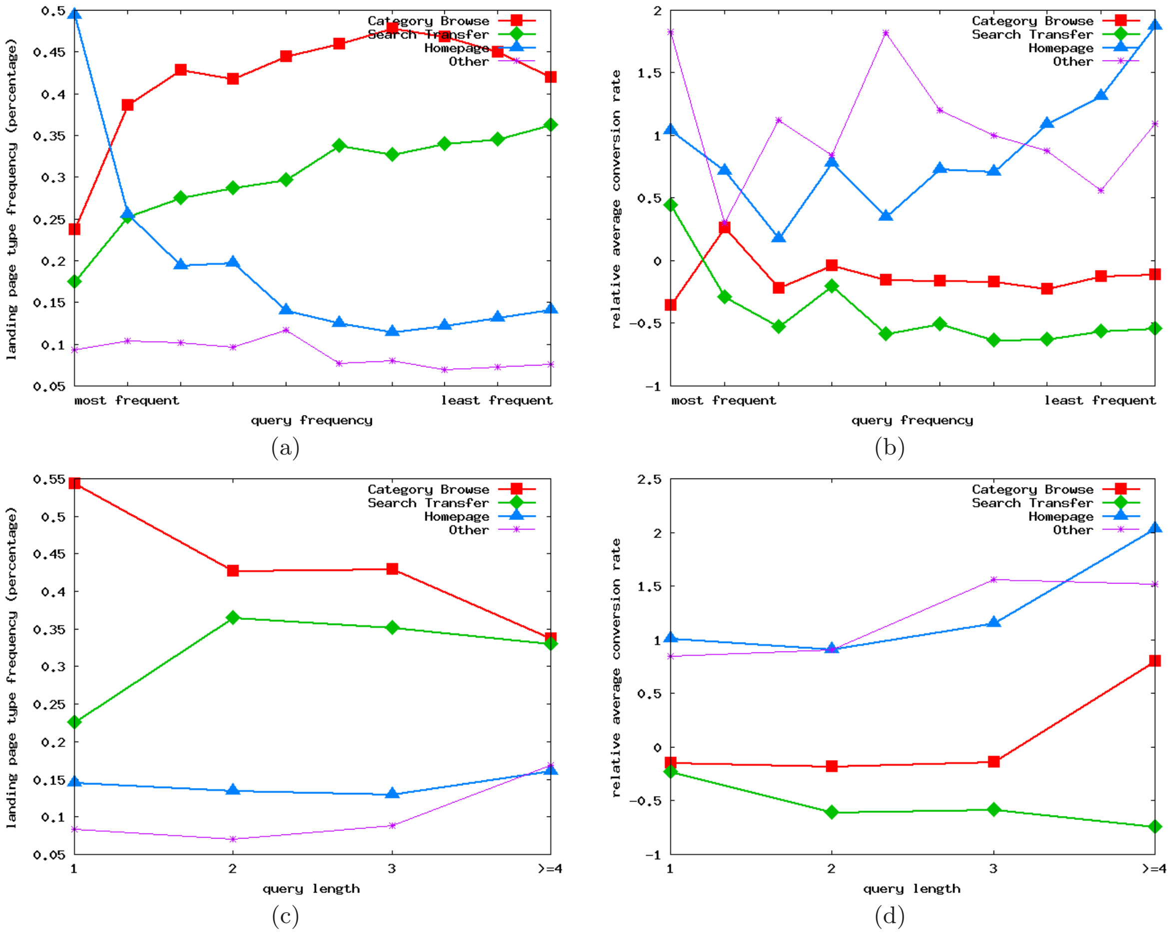

Figure 2: Landing page type for different types of queries on the opt-in dataset: which page types are popular choices (higher frequency in the left column) vs. which ones have higher conversion rates (higher relative average conversion rate in the right column)

3(d)). Still, a closer examination reveals a number of interesting details.

First, we observe from Figure 2(a) that Homepage is the dominating class used for the most frequent queries, and its usage gradually drops down as we move towards less frequent queries. Intuitively, the most frequent queries are more likely to be navigational queries or informational queries on popular brand names. Indeed, we examined the 100 most frequent queries in our conversion dataset, and found 43 of them to be brand names without any specific model indicators (e.g., nokia). In contrast, when less frequent queries included brand names, they tended to also include specific model information (e.g., 2009 chevrolet malibu). Note also that the usage of the Category browse and Search transfer classes gradually increase as we move to less frequent queries, with the usage of Category browse tipping off slightly towards the least frequent queries (reducing the gap with Search transfer). This indicates that as the queries become rarer, it is more difficult to pair them up with one of the pre-existing pages on the site (e.g., a Category browse page) and more convenient to resort to a Search transfer page. There is an interesting steady increase in the average conversion rate for the Homepage class as the queries become rarer (Figure 2(b)), in spite of it being the less popular choice. The conversion rates of the other two classes remain more or less constant for the 5 least frequent deciles.

Another characterization of query specificity is the length of the query. Longer queries are likely to be more specific (e.g., "100 polyester tablecloth" vs. "tablecloth"), although query length is not always a precise predictor of specificity (e.g., "asd2625kew[2]" vs. "christmas dinner recipe"). Note that the queries in the dataset do not cover a broad range of lengths, owing to the short average query length used in Web search today. Still, we observe that the difference between the usages of the Category browse and Search transfer classes are the widest for one-word queries, where the users are more likely to be looking for information at the category-level (Figure 2(c)). We also observe a similar increase in average conversion rate for the Homepage class as the queries get longer and more likely to be specific (Figure 2(d)).

Figures 3(a) and 3(b) present our analysis based on the price paid for the queries, which was used as a proxy for the query's commercial value since our conversion dataset does not contain auction information. The least expensive

[2]The intent of this query is most likely to be about a specific model: "Amana ASD2625KEW Side-by-Side Refrigerator."

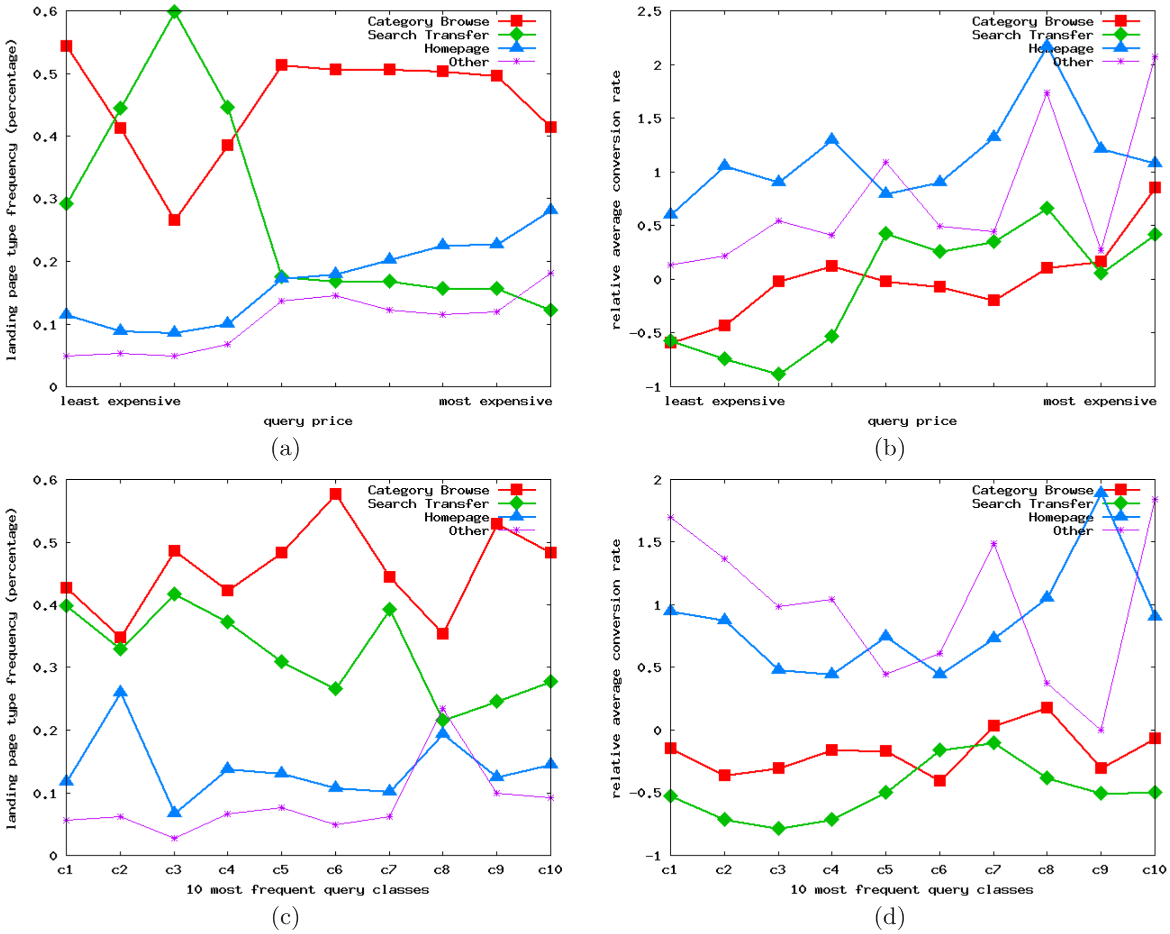

Figure 3: Landing page type for different types of queries on the opt-in dataset: which page types are popular choices (higher frequency in the left column) vs. which ones have higher conversion rates (higher relative average conversion rate in the right column)

queries are dominated by Search transfer and Category browse landing pages. As the queries become more expensive, there is a clear increase in the use of Homepage landing pages, as well as a drastic decline in the use of Search transfer landing pages. Interestingly, as the price goes higher, the general trends of average conversion rate for the three classes are all increasing overall. This suggests that the advertisers who paid more are not necessarily harder to "please." Indeed, these advertisers may be getting their money's worth as a result of higher quality landing pages, or better conversion of more expensive queries. While Search transfer pages have the lowest average conversion rate at the low price range, they yield higher average conversion rate than Category browse pages in the mid-price range. One possible explanation is that the low price range is dominated by low-quality Search transfer pages that are trying to monetize queries with lower commercial value, using less relevant landing pages or even spam or click arbitrage pages. Another possibility is that the low price range corresponds to less valuable keywords for which Search transfer provides a low effort solution.

In the next section, we further investigate the effectiveness of different landing page types on more comparable queries.

5.2.3 *Analysis of identical and related queries*

Here we examine different ad campaigns that targeted exactly the same queries. If advertisers used different landing page types for the same query, which type(s) tended to have higher conversion rates? Results are summarized in the left part of Table 4. It turned out that most queries were associated with only one landing page in this dataset, and conversions for multiple landing pages were reported for only around 600 queries. In order to obtain more reliable statistics, we relaxed the comparison to include different landing page types used for related queries, where two queries were considered related if they had at least one word in common and they shared the same query class (top one prediction from the query classifier). Results from this relaxed comparison study are reported in the right part of Table 4.

In both exact-match and relaxed-match studies, numbers reported in the i-th row and j-th column of each table encode two numbers $(w_{i,j} : l_{i,j})$, where $w_{i,j}$ denotes the number of times class i (c_i) out-numbers (out-performs) class j (c_j), and $l_{i,j}$ denotes the number of times c_i is out-numbered (out-performed) by c_j. $(w_{i,j} : l_{i,j})$ is shown in bold face when $w_{i,j} > l_{i,j}$. A class whose corresponding row contains many bold-faced entries tends to win in terms of either the click

On different landing pages used for the exact same query

(a) Click comparison:

	C. browse	S. transfer	Homepage	Other
C. browse	-	112:176	**72:50**	**33:31**
S. transfer	**176:112**	-	46:52	**21:17**
Homepage	50:72	**52:46**	-	**41:31**
Other	31:33	17:21	31:41	-

(b) Conversion rate comparison:

	C. browse	S. transfer	Homepage	Other
C. browse	-	17:57	**37:13**	**14:11**
S. transfer	**57:17**	-	**18:6**	**9:6**
Homepage	13:37	6:18	-	13:13
Other	11:14	6:9	13:13	-

On different landing pages used for related queries

(c) Click comparison:

	C. browse	S. transfer	Homepage	Other
C. browse	-	1514:2332	733:1046	422:752
S. transfer	**2332:1514**	-	**745:732**	379:523
Homepage	**1046:733**	732:745	-	338:460
Other	**752:422**	**523:379**	**460:338**	-

(d) Conversion rate comparison:

	C. browse	S. transfer	Homepage	Other
C. browse	-	263:824	**450:350**	259:278
S. transfer	**824:263**	-	**393:123**	**208:88**
Homepage	350:450	123:393	-	179:228
Other	**278:259**	88:208	**228:179**	-

Table 4: Comparison of different landing pages used for the same query or related queries. Comparison for both click and conversion rate between class i and j is summarized in the i-th row and j-th column by two numbers $w_{i,j} : l_{i,j}$, where $w_{i,j}$ is the number of times class i wins over class j (more clicks or higher conversion rate), $l_{i,j}$ is the number of times class i loses to class j. $w_{i,j} : l_{i,j}$ is shown in bold face when $w_{i,j} > l_{i,j}$.

competition (which class tends to get higher click) or the conversion rate competition (which class tends to get higher conversion rate). For instance, when landing pages from the Category browse and Search transfer classes were used for related queries, 2332 of the times the Search transfer page got more clicks, and 824 (263) of the times the Search transfer page got higher (lower) conversion rates. The numbers in Table 4 consistently reveal the Search transfer class to be much more likely to be the winner with higher conversion rate when compared against a page from another class used for either the same or related queries. This suggests that on fair comparisons Search transfer landing pages are quite effective at achieving conversions.

5.3 Browsing patterns as conversion events

When an advertiser uses a Homepage as landing page, presumably they are hoping to entice users to further explore the site via browsing. Compared to the other two dominant classes, the Homepage class is less likely to preserve the search context, especially for less common queries. Will the users be interested enough to continue browsing as expected or will they lose interest and leave the site immediately upon viewing the home page used as the landing page?

We use the afore-mentioned browsing dataset to answer this question. For each landing page in this dataset, the number of additional intra-site clicks in the same session can be extracted from the toolbar logs. If we define a click-based conversion as a visit where additional clicks on the same site exceed a threshold (three, in our case), we can then compute average conversion rate as defined in Section 5.1. As shown in Figure 4(b), overall we do observe the highest average conversion rate for the Homepage class. In fact, as the landing page gets more specific (Homepage → Category browse → Search transfer), additional clicks are less likely to occur. Clearly, one possible explanation is that upon landing on a page already very specific to the query, a user does not need as many clicks to arrive at a page that satisfies her. Still, our finding does show that even on rare queries, a more general-purpose landing page (e.g., a Homepage) does not defer users from further browsing.

Note also that while differing in details, the general trend of how the relative order of the three dominant landing page types, in terms of both usage (Figure 4(a)) and conversion (Figure 4(b)), changes across different query frequency remains consistent with our findings on the conversion dataset (Figure 2(a) and 2(b)). This demonstrates that our findings are not limited to one particular sample of advertisers represented in the conversion dataset.

6. CONCLUSIONS

In this paper we presented a study of context transfer in sponsored search advertising. By analyzing several hundred examples, we found that the majority of ad landing pages fall into three distinct classes *Homepage*, *Category browse*, and *Search transfer*. We then proceeded to build a machine learning classifier, capable of automatically mapping landing pages onto these classes. Using this classifier, we conducted a study of correlation between the different types of landing pages and the conversion rates of the corresponding ads.

We examined the suitability of different types of landing pages for different classes of queries by partitioning our data according to query frequency, length, topic, and price. We then studied the correlation of landing page types in each data partition with ad conversion rates. We analyzed several scenarios where choosing one type of landing page is preferable to the others. We also found that advertisers may favor landing page types that were not optimal for the queries that they were paired with.

Due to the variability in what constitutes conversion for different advertisers, in this paper we analyze correlation and not claim a causal relationship between landing page types and conversion rate. This limitation was introduced by the conversion data that was available to us. Nonetheless, this is a first attempt to provide insight into the relationship between landing page types, query classes, and conversion. For future work, we intend to study the causal relationship between landing page types and conversion, for groups of advertisers who measure conversion similarly. In addition, we plan to examine the correlation between conversion and other revealing data (e.g., query words, business category).

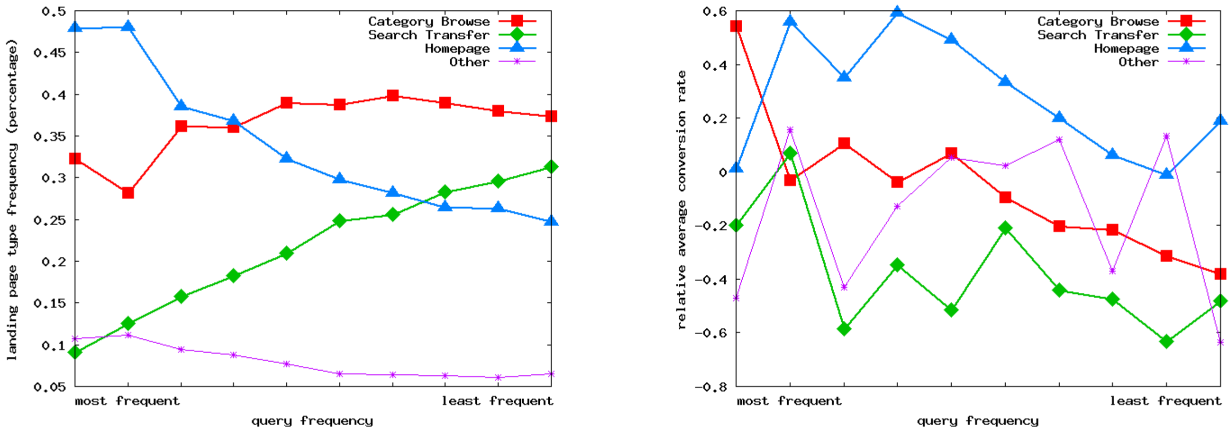

(a)
Percentage of different landing page types used by the advertisers

(b)
Relative average conversion rate

Figure 4: Additional browsing as conversions: study on the browsing dataset

Based on our findings, we encourage advertisers to experiment with different types of landing pages, and then make an informed choice based on statistical evidence.

7. REFERENCES

[1] H. Becker, A. Broder, E. Gabrilovich, V. Josifovski, and B. Pang. Context transfer in search advertising. In *SIGIR'09*. Poster.

[2] A. Broder. A taxonomy of web search. *SIGIR Forum*, 36, 2002.

[3] A. Broder, P. Ciccolo, M. Fontoura, E. Gabrilovich, V. Josifovski, and L. Riedel. Search advertising using Web relevance feedback. In *CIKM'08*, 2008.

[4] A. Broder, M. Fontoura, E. Gabrilovich, A. Joshi, V. Josifovski, and T. Zhang. Robust classification of rare queries using web knowledge. In *SIGIR'07*, 2007.

[5] D. Downey, S. Dumais, D. Liebling, and E. Horvitz. Understanding the relationship between searchers' queries and information goals. In *CIKM*, 2008.

[6] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[7] D. Fain and J. Pedersen. Sponsored search: A brief history. In *Second Workshop on Sponsored Search Auctions*, 2006.

[8] R. Jones and K. Klinkner. Beyond the session timeout: Automatic hierarchical segmentation of search topics in query logs. In *CIKM*, 2008.

[9] U. Lee, Z. Liu, and J. Cho. Automatic identification of user goals in web search. In *WWW*, 2005.

[10] Y. Li, Z. Zheng, and H. Dai. KDD CUP-2005 report: Facing a great challenge. In *SIGKDD Explorations*. 2005.

[11] P. Nelson. Advertising as information. *Journal of Political Economy*, 82(4):729–54, July/Aug. 1974.

[12] F. Radlinski, A. Broder, P. Ciccolo, E. Gabrilovich, V. Josifovski, and L. Riedel. Optimizing relevance and revenue in ad search: A query substitution approach. In *SIGIR'08*, 2008.

[13] M. Regelson and D. Fain. Predicting click-through rate using keyword clusters. In *Second Workshop on Sponsored Search Auctions*, 2006.

[14] B. Ribeiro-Neto, M. Cristo, P. B. Golgher, and E. S. de Moura. Impedance coupling in content-targeted advertising. In *SIGIR'05*, 2005.

[15] M. Richardson, E. Dominowska, and R. Ragno. Predicting clicks: Estimating the click-through rate for new ads. In *WWW'07*. ACM Press, 2007.

[16] D. Rose and D. Levinson. Understanding user goals in web search. In *WWW*, 2004.

[17] C. Wang, P. Zhang, R. Choi, and M. D. Eredita. Understanding consumers attitude toward advertising. In *8th Americas Conference on Information Systems*, 2002.

[18] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, 2 edition, 2005.