

WPI-CS-TR-12-03

August 2012

Understanding What They Do with What They Know

by

Craig E. Wills and Can Tatar

Computer Science
Technical Report
Series



WORCESTER POLYTECHNIC INSTITUTE

Computer Science Department
100 Institute Road, Worcester, Massachusetts 01609-2280

Abstract

This work seeks to understand what “they” (Web advertisers) actually do with the information available to them. We analyze the ads shown to users during controlled browsing as well as examine the inferred demographics and interests shown in Ad Preference Managers provided by advertisers.

In an initial study of ad networks and a focused study of the Google ad network, we found many expected contextual, behavioral and location-based ads along with combinations of these types of ads. We also observed profile-based ads. Most behavioral ads were shown as categories in the Ad Preference Manager (APM) of the ad network, but we found unexpected cases where the interests were not visible in the APM. We also found unexpected behavior for the Google ad network in that non-contextual ads were shown related to induced sensitive topics regarding sexual orientation, health and financial matters.

In a smaller study of Facebook, we did not find clear evidence that a user’s browsing behavior on non-Facebook sites influences the ads shown to the user on Facebook, but we did observe such influence when the Facebook Like button is used to express interest in content. We did observe Facebook ads appearing to target users for sensitive interests with some ads even asserting such sensitive information, which appears to be a violation of Facebook’s stated policy.

1 Introduction

There has been much work measuring the extent to which third-party advertisers are in a position to observe and correlate user behavior across a broad range of first-party Web sites [20, 2, 16]. Other work has shown that these advertisers are not only in a position to infer information about users, but directly obtain known information about users via social networking and other types of sites [15, 17, 14]. Recent work summarizes the policy and technology issues on third-party Web tracking [18], but an important question that has gone largely unanswered, which we seek to investigate in this work, is understanding what “they” (the advertisers) actually do with this information available to them.

We are aware of two previous studies that are relevant to this question. The first gathered information on text ads shown to automated users visiting different first-party sites [13]. The researchers were able to determine that different ads were shown based on different sites visited, but they did not examine the nature of these differences. The second performed a controlled study to measure behavioral targeting for four training topics on text ads as a basis for evaluating the effectiveness of privacy tools designed to limit this targeting [4].

Our work makes a number of contributions. First, we not only examine how advertisers use behavioral information in serving ads, but take a more comprehensive approach to see if and how this information is combined with location and personal characteristics of a user. Second, we introduce a variety of sensitive topics in our testing to understand how they are handled by advertisers. Third, we examine not just text ads, but also those ads shown as images or Flash objects.

In addition to what ad networks reveal about what they know via the ads they display, some have provided transparency for users to see what inferred demographics and interests are being associated with their browsing behavior. These Ad Preference Managers (APMs) allow users to view and edit their preferences. Another contribution of our work is to examine the behavior of these APMs as they provide additional insight on what advertisers are doing with what information they receive. We are not aware of any previous work that has examined APMs, let alone their completeness and accuracy.

The key idea of our approach is to control the set of inputs that an ad network receives and use the output of the ad network—the contents of the APM and its displayed ads—to understand what is being done with this information. We use a modification of our methodology in a less-extensive study of Facebook to study how it translates user information and a user’s behavior on non-Facebook sites into ads served on the social networking site. Recent work such as [21], has demonstrated that Facebook is in a position to track user behavior on non-Facebook sites.

The remainder of the paper begins with an overview of APMs in Section 2. We then review what information is received by advertisers and present a classification for how it is used to serve ads in Section 3. Section 4 describes our testing methodology and Section 5 summarizes initial results. We go on to describe a more extensive study and results for the Google Ad Network in Section 6 followed by a smaller study of Facebook in Section 7. We conclude with a summary of our findings and directions for future work. An appendix of illustrative examples captured during testing (and referenced in the paper) is provided for the interested reader.

2 Ad Preference Managers

As a response to be more transparent on how user browsing behavior is translated into preferences for advertising, a number of third parties provide interfaces for users to view and edit their preferences. The BlueKai Registry [5] was one of the first such managers. In 2009 Google announced its Ads Preferences Manager [10]. In our preliminary work we identified approximately ten ad networks that provide an “Ad Preference Manager” or APM. Figure 1 shows an example of an actual Google APM we captured during our work where user behavior has been translated into four interest and two demographic categories.



Figure 1: Google Ads Preferences Manager

How browsing behavior is translated into APM topics and ultimately into displayed ads is particular to each ad network. However we establish expectations for behavior based on the example in [11] where an Internet user “Mary” has a favorite hobby of gardening and visits many gardening Web sites. It is expected behavior that an ad network will include the topic of “gardening enthusiast” in her APM and serve related ads. It would be unexpected (or at least unexplained) if Mary’s APM included the topic of “sports”, but Mary did not visit any such sites. It may also be the case that Mary sporadically visits other types of sites such as those for automobiles in looking to buy a car. While the topic of “automobiles” could be shown in the APM, the ad network may not be certain that this topic is of interest and decide not to show it in the APM nor use it in behavioral advertisements. Sensitive topics are typically not included in APMs with policies such as Google’s explicitly stating as much [11].

3 Ad Classification

Previous work has shown that third-party advertisers receive much information about user browsing behavior across an increasing number of Web sites [20, 2, 16]. These advertisers not

only learn which Web pages are visited by users via the HTTP `Referer` header, but also in many cases receive text input when users enter a search query on a page. Advertisers use the contents of the current page to display *contextual* ads on the page and *infer* characteristics about the user for use in displaying *behavioral* ads as the user visits other Web pages.

In addition to inferred information, other work has shown that third-party advertisers receive *known* information about a user via sites in which the user creates an account and provides profile information to the site. For example, early work found such leakage via online social networking sites [15] while more recent work has found similar leakage for many other types of sites [17, 14]. One example of such leakage occurs when a user’s music interests, age, gender and zip code from Pandora are sent to DoubleClick (a Google property) via an HTTP GET request (shown in Figure 8 of the appendix). Similar leakage is done via LinkedIn, which leaks information such as job title, job function, company, zip, college, graduation year, gender and age from a user’s LinkedIn profile to DoubleClick.

As a means to help understand what third parties do with the information they obtain we have developed a two-dimensional classification with one dimension for known information and one dimension for inferred information. This classification is shown in Figure 2 and focuses on the ads served by an ad network.

The horizontal dimension shows information that is known about a user where a user’s current location, obtained via a precise latitude/longitude location or a geoIP mapped location, is one type of known information about a user. We note that location could be shown as a third dimension, but add it to the known information dimension for simplicity of display. The resulting horizontal dimension indicates ads may be displayed using no known information about a user, using the user’s current location, using information from a profile available on the current page (e.g. using information from a user’s Pandora profile to show ads while the user visits Pandora) and using information obtained from a previous profile to show an ad on a subsequent page (e.g. using information from a user’s Pandora profile for displaying an ad on a subsequent page).

The vertical dimension represents information inferred about a user’s browsing behavior including no behavior, context of the current Web page and past browsing behavior. As shown in the figure, we separate past browsing behavior information into three types: behavior captured in the Ad Preference Manager for the user, behavior not captured in the APM and behavior that deals with sensitive topics.

The boxes within the figure show potential combinations of information used for an ad where we adopt the notion of a “generic ad” for an ad that has known no context on either dimension. While the CAPITALIZED TEXT in the figure labels ads based on only known or inferred information we expect combinations also exist—such as ads combining context of the current page along with a user’s current location.

The shading of the figure groups combinations into *expected* and *unexpected* based on what we believe would be the perception of most users. Thus on the horizontal dimension we expect to see location and profile information for the current page being used while on the vertical dimension we expect to see contextual advertising as well as behavioral advertising that is captured in the APM of the associated ad network. We do not expect to see profile information being used on subsequent pages nor do we expect to see behavioral ads for topics not captured in an APM or such ads dealing with sensitive topics.

We use the classification and its expected and unexpected ad behavior to drive the

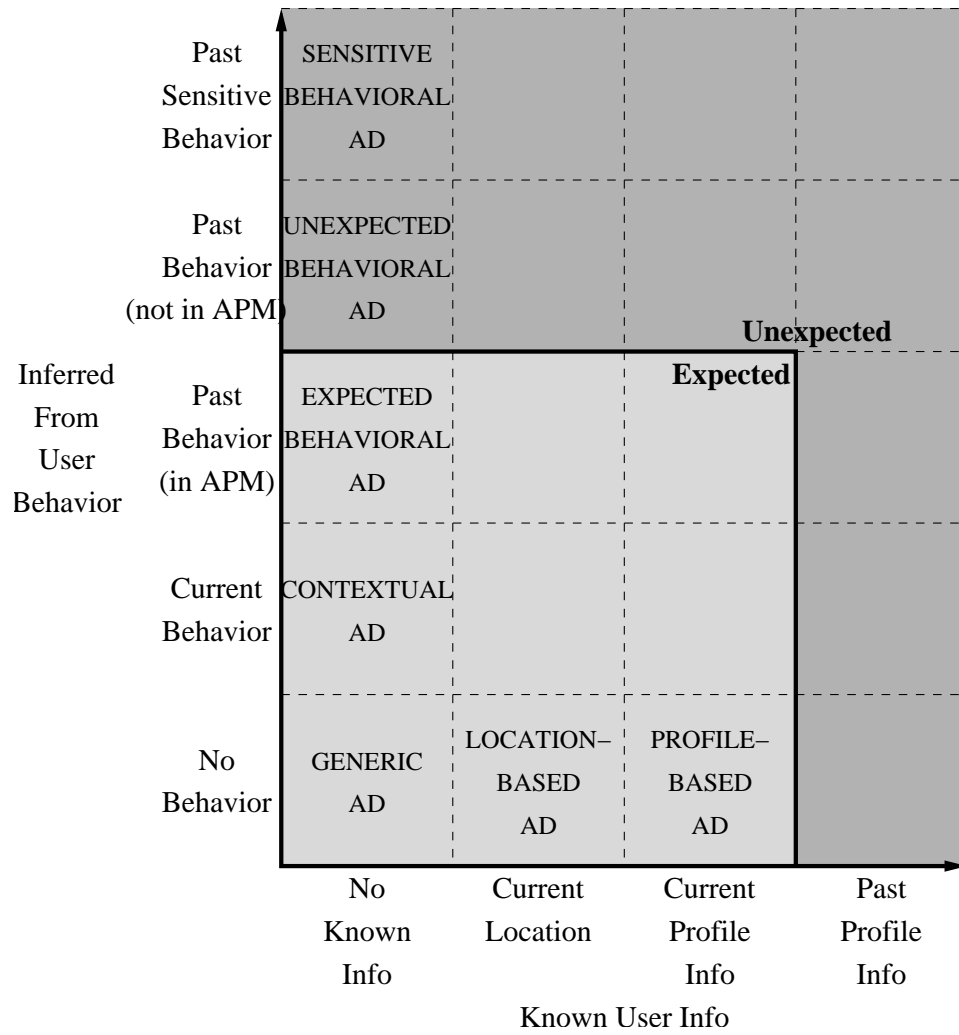


Figure 2: Classification of Ads Shown

research questions that we seek to answer in our research study. These questions are:

1. What expected combinations do we see being used in the serving of ads?
2. Do we see evidence of profile information obtained on one site being used to serve ads on subsequent sites?
3. Do we see evidence of ads shown matching a user’s behavior, but this behavior not being reflected in the APM?
4. Do we see evidence of ads shown based on previously-viewed sensitive topics?

4 Methodology

Our approach for answering these questions is to study specific ad networks. An ad network is largely a black box in terms of how it works. However, an ad network does have a number of inputs that can be controlled and outputs that can be observed as shown in Figure 3. Specifically the sites that are visited and the textual input (search terms) provided to an ad network can be controlled. Similarly, current location and profile information can be controlled. Manipulating these inputs in a controlled manner allows us to detect if they come through as visible output—in the APM of the ad network, in the displayed ads themselves or in both.

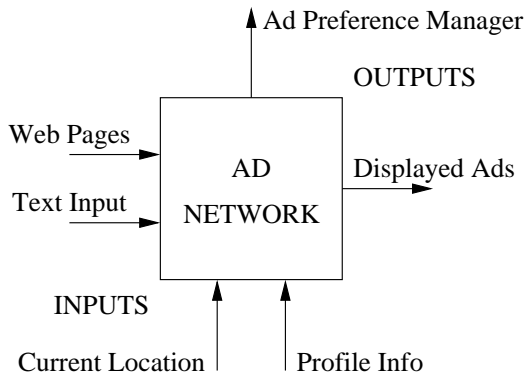


Figure 3: Inputs and Outputs of an Ad Network

Our first step in studying a particular third-party ad network is to identify a set of first-party sites for which the advertiser has a presence when the first-party site is loaded in a browser. We primarily found these first-party sites by using results from an ongoing longitudinal Web crawl using the methodology described in [16]. We also include sites dealing with sensitive topics (e.g. health and sexual orientation) as well as sites where a user creates an online profile of information. We identify about 15 first-party sites for each ad network.

Each experiment for an ad network consists of daily sessions for a ten-day period. A session comprises visiting the first-party Web sites successively and performing representative actions on these sites (such as viewing the technology page on a news site or searching for information on skin cancer on a health site).

The same controlled browser is used for an entire experiment. At the beginning of the experiment we delete all cookies and history in the test browser. We do not explicitly delete Flash LSO cookies, but did not observe any such cookies for the ad networks we tested. For sites requiring a login, we do so in the first session of the experiment and then do not log out during the remainder of the experiment. All tests are run from Worcester, Massachusetts, which is part of the Boston metropolitan area.

During an experiment, we observe the ads served by the relevant ad server and check the APM contents during the course of the sessions to see how it changes based upon the visited sites and inputs. We do not click on any ads during a session thus providing no additional information about interests. We also record all HTTP traffic and object content using the Fiddler proxy [9] running on the same computer as the test browser. This traffic and content is saved for later analysis.

For this work, all tests are performed manually with a session typically taking 20-30min. We plan to automate the tests moving forward, but the manual approach not only allows us to drive the test, but to observe and understand the rendered ads whether they be text, image or Flash. These observations helped us in automating the analysis of the collected content. We are also able to stop and capture examples of different types of ads based upon our classification in Figure 2. Because the tests are performed sequentially, the results are impacted by the churn of ads over time. Because of this churn we are cautious in attaching significance to small differences when reporting results from separate experiments.

5 Initial Results

We only applied our methodology to ad networks that provide an APM. From this set of ad networks, we initially studied four of the larger ones: AOL [3], BlueKai [5] (actually a data exchange), Google [10] and Yahoo! [22].

We performed an experiment for each of these four networks and observed the resulting APMs and rendered ads. For AOL and Yahoo!, we observed evidence of generic, contextual, location and behavioral advertising consistent with what is shown in the APM. Each of these types of ads is expected based upon our classification in Figure 2. We observed these same types of ads for Google, but also observed behavioral ads for topics not shown in the APM and what appeared to be behavioral ads for sensitive topics. We also observed profile-based advertising. BlueKai does not directly serve ads, but we did observe HTTP redirections to a number of ad servers. However, we were unable to characterize the types of ads served based on information known by BlueKai.

Based upon these initial results, we choose to focus on more systematic analysis of the Google ad network for a number of reasons: it is the largest ad network with a presence on approximately 60% of popular Web sites based on our own recent data collection using the methodology described in [16]; it was the only ad network for which we saw evidence of unexpected results from our initial work; and on March 1, 2012 it modified its privacy policy [12] in part to show “more relevant ads” affording an opportunity to study how activities on Google-owned sites affects ads on non-Google sites.

6 Google Results

With this focus, we first provide additional details about the experiments performed in testing the Google Ad Network then present the results we obtained from the collected data.

6.1 Google First-Party Sites

Table 1 shows the set of 20 first-party Web sites used to test the Google Ad Network. The set of sites was increased from the initial set to add new induced interests, reinforce others and add sites for the display of ads. The category shown for each site is based upon characterizations provided by the Web traffic analysis companies Alexa [1], Compete [6] and Quantcast [19]. We note that this set of sites includes no Google-owned first-party sites. Sites were used to search for a topic and browsed by visiting top-level tabs. On the dictionary sites, we provided random strings of characters that were non-English words. This approach yielded contextual ads for improving English skills, but otherwise resulted in ads without context on the page.

Table 1: Sites Used for Google Ad Network

Web Site	Category
bloomberg.com	Financial News
accuweather.com	News/Weather
tripadvisor.com	Travel Planning
yelp.com	User Reviews
ford.com	Automotive
toyota.com	Automotive
gaylife.about.com	Gay Life
thenewgay.net	Gay/Lesbian
linkedin.com	Professional Networking
pandora.com	Radio
medhelp.org	Health/Support
menshealth.com	Men/Health
metrolyrics.com	Music/Lyrics
tmz.com	Entertainment
cbsnews.com	News
cnn.com	News
huffingtonpost.com	News
nytimes.com	News
macmillandictionary.com	Dictionary
thefreedictionary.com	Dictionary

Using this set of sites, we performed six experiments with a total of 55 sessions (including one session of 8 days and one of 7) in May/June 2012 where in each experiment we induced different sets of interests. Details on the interests and how they were induced are described in

Section 6.4. An Internet Explorer browser with a default configuration (all cookies accepted, JavaScript execution enabled, no ad blocking) was used for all experiments.

The saved content from each session was analyzed in an automated manner where we only considered content that was served directly or indirectly by the Google ad network itself. Indirection was detected via the string “google” or “doubleclick” being in the `Referer` header or as part of JavaScript code that was serving the ad. We used a keyword-based analysis for each induced interest where we searched the saved content for one of a set of keywords related to that interest. Automated analysis of image and Flash content was done by matching keywords in URLs. Matches were subsequently verified to ensure correctness and that the ad was not contextual based upon the contents of the page.

6.2 Generic, Contextual and Location Ads

We use the classification in Figure 2 to present our results. Starting with the four combinations in the lower left corner of the figure, it is not surprising that we observed numerous ads each session related to the contents of the current page (contextual advertising) or having no clear relationship to the current or previous behavior (generic advertising). A representative example of a contextual ad shows scientific brain games on a `nytimes.com` Science page (Figure 9 in the appendix).

Location-based advertising is also common. Analyzing the data for the keywords “worchester,” “boston,” “massachusetts” and “new england,” we verified that all 55 experimental sessions exhibited at least one instance of location-based advertising, such as shown in Figure 10.

It is also not unusual to see the combination of contextual and location-based advertising such as a `match.com` ad appearing on `gaylife.about.com` (shown in Figure 11). The ad is both location-based and contextual by inviting the user to “View Gay Men in Worcester” and showing photos of men.

6.3 Profile-Based Ads

We next looked for evidence of profile-based ads (two right-most columns in Figure 2) using the sites `linkedin.com` and `pandora.com` in which we established an account with the profile information described in Section 3. Although `linkedin.com` transmits much profile information to DoubleClick, we did not find evidence that this information was being used by DoubleClick in serving ads on `linkedin.com`. We did observe `linkedin.com` serving many ads on its own site (and other sites), but did not explicitly study these ads.

However, we did find evidence that profile information on Pandora is being used when DoubleClick serves ads on the site. We found two types of ads (Figure 12) where the ads match information in the Pandora profile. The first matches the profile location of New York while the second shows an ad from `match.com` with a default age range that corresponds to the age in the profile. It is interesting that this ad makes use of the current location of Worcester rather than profile location of New York. This example also shows how two different `match.com` ads were each served using the current location (Worcester), but this ad uses the Pandora profile to show a particular age range and photos of women while the previous contextual ad on `gaylife.about.com` shows photos of men.

In Figure 2 we show profile-based advertising as expected behavior because in this case DoubleClick and Pandora are partners. However, we would not expect the profile information to be used for advertising on subsequent first-party sites, which is this last column in Figure 2. Evidence for such subsequent use is not clear. We found no evidence that the age in the profile is being used as the Google APM consistently converged on demographics of a 65+ age male in all of our experiments such as shown in Figure 1. The profile age did not influence this inferred characteristic nor did we observe the use of the profile age in ads except those for `match.com` on Pandora.

Evidence for the subsequent use of the profile location is less clear. Using our automated analysis with keywords “nyc” and “new york,” we examined all sessions where the profile location of New York was passed to Google via LinkedIn and Pandora. We found that 59% of such sessions contained at least one ad (not on LinkedIn or Pandora) for this location with typical examples being a NYC campaign for Liberty Medical shown on `medhelp.org` or ads for New York hotels on `gaylife.about.com`. Given that New York is a popular location in general, it is expected that some number of generic ads for it would appear. Unfortunately, we did not perform any experiments without LinkedIn and Pandora so we do not have a baseline for comparison.

6.4 Induced Behavioral Interests

We next test for the presence of behavioral advertising (rows 3 and 4 in Figure 2) by inducing a number of interests in each of our experiments. The complete set of interests, along with how they were induced and the keywords used when we analyzed the collected data are shown in Table 2.

Table 2: Induced Behavioral Interests

Induced Interest	How Induced?	Match Keyword(s)
cars	Ford, Toyota sites	ford, toyota, cars, autos, mazda, honda, jeep
dogs	search term	dog, k-9, pets, veterinarian, puppies
golf	search term	golf
investment	Bloomberg site	finance, invest, stock, market, trusts
miami	location selection	miami, south beach
tennis	search term	tennis, racquet

As shown in the table, two interests were induced by visiting relevant Web sites—`ford.com` and `toyota.com` for an interest in buying a car and `bloomberg.com` for an interest in financial investments. Three other interests—dogs, golf and tennis—were induced as a search term on one of the news or dictionary sites. Previous work [4] had induced interests only by visiting relevant Web sites, but in preliminary work we found that text input on a site could also be used to induce an interest. Finally, a destination location of interest, in this case Miami, was induced by selecting/entering it on `accuweather.com`, `tripadvisor.com` and `yelp.com`. In all cases, the Google Ad Network learned of these interests via HTTP requests to `doubleclick.net`. Within an experiment, more than one, but fewer than all interests were induced. Inducing multiple interests within the same experiment is in contrast

to the methodology of [4] where only one was induced at a time, but we believe that our approach represents more realistic browsing behavior.

6.5 Evolution of Ad Preference Manager

Before examining results on how these induced interests are mapped into behavioral ads, we first examine how these interests influence the evolution of the Ad Preference Manager—the other output of the ad network—over time. Figure 4 shows the evolution of the APM over eight sessions (days) of one experiment where cars, golf, investments and miami were the induced interests from Table 2.

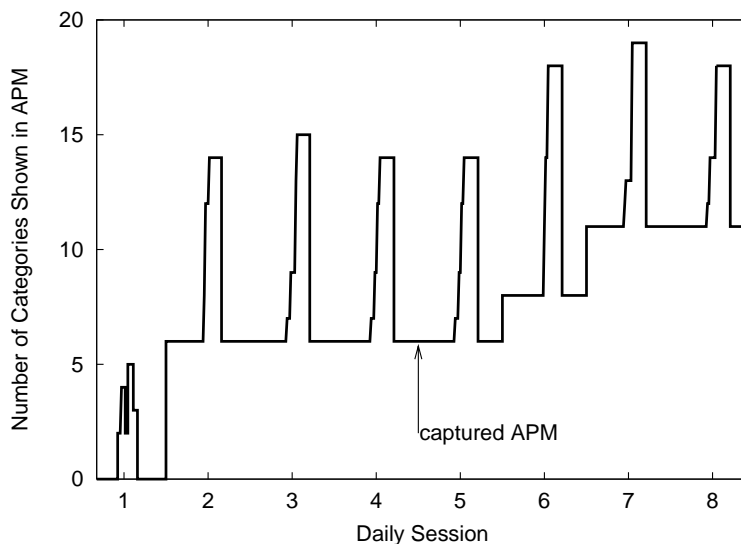


Figure 4: Evolution of APM Over Time

Figure 4 shows the number of interest and demographic categories displayed by the Google Ads Preferences Manager [10] at the beginning of, during and after each session. The horizontal gaps between each session represent elapsed time in which no Web site is visited with our test browser. At no point during the experiment were categories manually added or removed; what is shown reflects only inferences made based upon browsing behavior. Figure 1 shows the contents of the APM as it was captured at the indicated point in Figure 4. At that point in time, the APM contained the four interest categories and two demographic categories for the total of six categories shown in Figure 1.

While the process used by Google for adding and removing APM categories is not known, the behavior in Figure 4 appears to reflect a two-stage process for mapping browsing behavior to the set of categories. We see short-term behavior where interest (not demographic) categories are added in response to input text and the content of particular pages. Observation shows that visiting pages of `nytimes.com` and to a lesser extent `bloomberg.com` results in immediate addition of categories related to the content of these pages. In contrast, we do not observe this behavior when other news sites such as `cnn.com` or `cbsnews.com` are visited. However these short-term additions are fleeting as we observe each of them being

dropped from the APM within an hour of when the session is completed resulting in the spikes shown in Figure 4.

The addition of categories that persist long-term occurs between sessions. For example in Figure 4, we see that the number of categories in the APM returns to zero shortly after the completion of the first session, but by the time the second session is started the next day, six (four interest, two demographic) categories have been added and these categories persist across the short-term comings and goings of categories in subsequent sessions. These longer-term categories tend to change little based upon our observations. The four categories shown in Figure 1 correctly reflect the induced interests of golf and investments as well as interest in the weather based on the site `accuweather.com`. Subsequently added long-term categories in Figure 4 do not match the induced interests, although the category for “Miami” does finally appear for the first time as a category after Session 8. The Google category of “Autos & Vehicles” or categories specific to Ford or Toyota did not appear at all.

6.6 Behavioral Ads

With a better understanding of how the APM works, we now examine the other visible output of the Google Ad Network—the ads themselves. Each of the interests in Table 2 was induced for some number of the six experiments as indicated. In other experiments, these same interests were explicitly not induced; for example by not entering it as a search term in the case of dogs or by not visiting `ford.com` and `toyota.com` in the case of cars.

In analyzing the results we use the keywords (and appropriate variants) shown in Table 2 to match ads in each session that were served by Google and are not contextual—for example ads for financial investments on `bloomberg.com` are contextual as are ads for golf on a news page containing an article about a golf tournament. Behavioral ads are non-contextual ads corresponding to an induced interest. In our testing we analyze the number of non-contextual ads for an interest both when it is induced and not induced. Two representative behavioral ads we observed were for the induced interests of golf and dogs (shown in Figure 13) where in each of these cases, the corresponding categories of “Sports - Individual Sports - Golf” and “Pets & Animals - Pets - Dogs” appeared in the APM.

These behavioral ads for interests that appear in the APM are expected as classified in Figure 2. We also observe combinations where behavior is combined with location for ads such as “Puppies for Sale in Massachusetts” on `thefreedictionary.com`.

Figure 5 shows a summary of results for the six interests given in Table 2. In each case, the results show the percentage of sessions in which at least one non-contextual ad match occurs for the given interest. Results are provided for sessions in which the interest is not induced and sessions for which it is. Note the interest golf was induced in all sessions so non-induced results are not-applicable in the figure. All results are based on at least 10 sessions with most results based on more than 20 sessions.

Figure 5 shows a number of interesting results. Non-contextual ads for dogs and tennis were only found when the interest was induced and in each case corresponding categories populate the APM. In contrast, ads for cars and financial investment always or almost always are found whether or not the interest is explicitly induced. What is unexpected about this result is how the Google APM is handled in each case. For investment, the Google category “Finance - Investing” appears in all relevant experiments while for cars, we never saw a

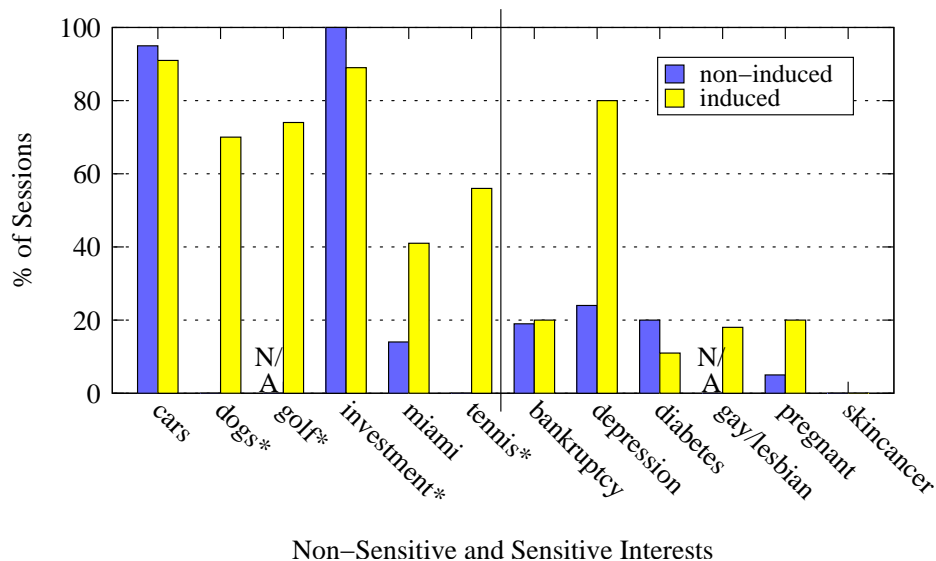


Figure 5: Percentage of Sessions Displaying Non-Contextual Ad Matching Induced Interest (* indicates interest was shown in APM when induced.)

category for automotives in general or Ford or Toyota in particular appear in the APM during an experiment.

We observed similar behavior for the interest Miami. We saw a number of non-contextual ads for Miami hotels (such as shown in Figure 14) when the interest was induced (and a lesser number when it was not). However in most of the experiments where it was induced, the interest was not reflected in the APM. In only one experiment did we see the category of Miami eventually reflected in the APM as described for Figure 4.

It can be argued that the absence of an expected interest in the APM is not an issue. What is shown in the APM is just a heuristic and a user can add the interest themselves. However if the APM is accurate for some interests and not for others then it is inconsistent and not complete in representing what is known about the user, which was part of the rationale in introducing APMs in the first place.

6.7 Induced Sensitive Interests

The top row of Figure 2 shows that behavioral ads could conceivably be served for sensitive topics. We next repeated our methodology for a range of sensitive interests to test for evidence of such ads in light of the policy that “Google will not show ads based on sensitive information or interest categories, such as those based on race, religion, sexual orientation, health, or sensitive financial categories” [11]. Consistent with this statement, we observed that induced sensitive interests had no effect on the APM categories.

The interests we tested included sexual orientation as well as ones on sensitive health and financial matters. These interests are shown in Table 3 along with how they were induced and the keywords used when we analyzed the collected data. We use the `gaylife.about.com` and `thenewgay.net` sites to induce the gay/lesbian interest, the two sites `medhelp.org` and `menshealth.com` to induce the three health-related interests and one of the news or

dictionary sites to induce bankruptcy.

Table 3: Induced Sensitive Interests

Induced Interest	How Induced?	Match Keyword(s)
bankruptcy	search term	bankrupt, chapter 7, debt, tax relief, foreclosure
depression	health search term	depression
diabetes	health search term	diabetes
gay/lesbian	gaylife, thenewgay sites	lgbt, lesbian, gay, mat_boy
pregnancy	health search term	pregnant, ob/gyn, infant, baby, birth
skin cancer	health search term	skin cancer, melanoma

Despite Google’s statement, we do see non-contextual ads for induced sensitive interests in our collected data. One example (shown in Figure 15) is where the topic of depression appears in a group of three ads on `nytimes.com`. Interestingly, another ad in the group is for the Volkswagen Golf, a car model, where the Volkswagen brand did appear as a category in the APM for a short term in response to a search for “golf” earlier within the session in which this ad occurred.

Similarly, we found ads (such as Figures 16 and 17) that match the induced interests of bankruptcy and pregnancy. One of the ads is also contextual as it appeared on the `accuweather.com` page for Miami weather, which was being induced as an interest. We also observed an ad for a New England center for “Getting Pregnant After 30” on `macmillandictionary.com` combining a sensitive induced topic along with location. Similarly we observed instances of `match.com` ads with photos of men (similar to Figure 11) on non-gay sites as well as ads advocating “LGBT for Obama” on `thefreedictionary.com`.

These examples clearly show that the Google Ad Network is serving non-contextual ads related to induced sensitive topics. One obvious question is what behavior is observed when the topics are not induced. These results are shown in the right portion of Figure ???. Note: gay/lesbian was always an induced interest and we did not observe any non-contextual ads related to skin cancer. Again all results are based on at least 10 sessions with most results based on more than 20 sessions.

The overall results reflect that non-contextual ads for sensitive interests are shown less frequently than those for non-sensitive interests in Figure 5. The results also reflect the frequency that such ads are shown does not differ significantly except for depression where 8 out of the 10 (80%) sessions in the one experiment in which it was induced contained an ad on `nytimes.com` (similar to the one shown in Figure 15). We initially suspected these would actually be an example of contextual advertising, but found no evidence of “depression” in the contents of the `nytimes.com` page and have no basis to conclude they are contextual. We did ascertain many other ads for sensitive topics as being contextual and are not shown in Figure ???. For example, we found many ads related to bankruptcy on `cbsnews.com` at a period of time that the “debt crisis” was a featured topic for this site. We also did not include results if the contextuality of an ad was questionable. For example, pregnant results in Figure ??? do not include sessions in which ads for “baby names” appeared on dictionary sites as one could construe a list of names as similar to a list of words.

In summary, Google may not be serving ads based on sensitive interest categories, but it is serving non-contextual ads related to sensitive interests—whether induced or not. We do not know why, but one possible explanation is that Google includes ads for sensitive topics in its pool of ads that are served in non-contextual situations. The result is that these ads are served independent of whether or not a user’s behavior indicates a sensitive interest. Users with a sensitive interest are then served ads that may appear to them to be behavioral.

6.8 Induced Interests from Google Sites

In the last portion of our study, we examined some of the impact of Google’s modified March 1, 2012 approach to sharing information across Google properties [12]. We used this opportunity to study how activities on Google-owned sites affects ads on non-Google sites. For this part of our study, we changed the source from which interests were induced. Specifically, in some experiments the interests of golf, bankruptcy and depression were induced based on queries to the Google search engine. Similarly, YouTube was used to search for and view videos based on the interests of dogs and pregnancy. Finally in two experiments we performed them while logged into a Google+ account where the profile location was New York.

We found some interesting results from this portion of the study. The profile location portion was unfortunately confounded by the same location as LinkedIn and Pandora, but we did observe in one experiment using the Google+ profile that a New York City zip code was being consistently passed directly to `accuweather.com` demonstrating an integration between information in this profile and the Google Ad Network. We did find such sharing in experiments when logged out of Google+.

We also observed that interests induced via YouTube were reflected in ads and in APM categories. For example, we observed two non-contextual ads (shown in Figure 18) that were served on `thefreedictionary.com` within a session in which dogs and pregnancy were induced on YouTube. The inducement of dogs caused this category to be included in the Google APM in a similar manner as if it was induced from a non-Google site. In contrast, we did not observe that the inducement of golf as a Google search term caused this topic to be included as an APM category.

Figure 6 shows the impact of inducing sensitive and non-sensitive interests via a Google property. Results for these interests are repeated from Figure ?? and combined with results when the interest is induced from a Google property. All Google-induced results are for more than 20 sessions. The results show that induced interests using Google search show similar frequency as when the interests are not induced at all while the two YouTube-induced interests correlate better in frequency with induced interests on non-Google sites.

7 Facebook Results

We also used the ads served by Facebook to its users as a means to understand how the social networking company makes use of the information received from its users. With a single Web site to collect information and serve ads, it is an easier ecosystem to study. This

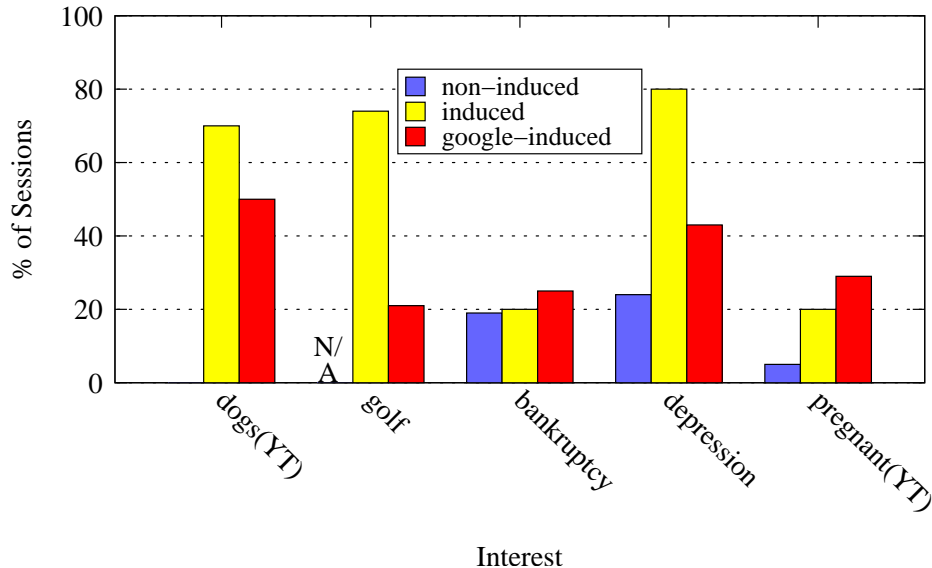


Figure 6: Percentage of Sessions Displaying Non-Contextual Ad Matching Interest for Interests Induced via Google Properties (YT indicates YouTube)

work was partially motivated by studies such as [21] and our own work showing Facebook as a growing *third-party* with a presence on nearly 40% of popular first-party sites.

This portion of our study was less extensive than our previous study and examined two specific research questions. First, in comparison with information provided by a user on Facebook itself how does a user’s browsing behavior on non-Facebook sites influence the ads this user sees on Facebook both when the content is and is not Facebook-liked? Second, how is sensitive browsing behavior and information about a user handled in ads by Facebook?

7.1 Influence of Non-Facebook Browsing Behavior

To answer these questions we used a similar methodology to that described in Section 4 in that we constructed a set of first-party sites having the Facebook Like button that was being served by Facebook. We did not include sites where Facebook graphics are provided by a first-party server so that a Facebook server is only contacted if a user clicks on the button. The first-party sites induced both non-sensitive and sensitive interests.

Using these sites and interests, we conducted four experiments, each of 15-20 sessions, where at the beginning of each experiment we logged into Facebook and remained logged in for the remainder of the experiment. At the end of each session we visited Facebook then saved the results of our browsing session for later analysis. In the first experiment, we did not induce any interests, but simply recorded the displayed Facebook ads in each session. Second, for each session we visited the Web sites in our test set, but did not click on the Facebook button. This experiment examines whether non-Facebook behavior is influencing Facebook ads. Third, we repeated the previous experiment, but in this experiment we not only visited sites, but Facebook-liked pages we browsed on these sites. Fourth, we did not visit any non-Facebook sites, but instead induced interests by indicating them as Facebook interests at the beginning of the experiment. We also indicated that our male user was

interested in men.

Results for non-sensitive and sensitive interests from these four experiments are shown in Figure 7. As shown, Facebook ads rarely match non-induced or non-Facebook-behavior-induced interests. The only instances for the latter case are ads for a Facebook fishing game and ones about college debt, which may relate to bankruptcy. In contrast, Figure 7 shows generally a high percentage of sessions included ads for interests induced either by Facebook-liking a page or explicitly including the interest as a Facebook interest. Example ads include one for biking in the New York City area (such as shown in Figure 19). Collectively, these results do not provide clear evidence that a user’s browsing behavior on non-Facebook sites influences the ads shown on Facebook itself unless a user Facebook-likes the site content.

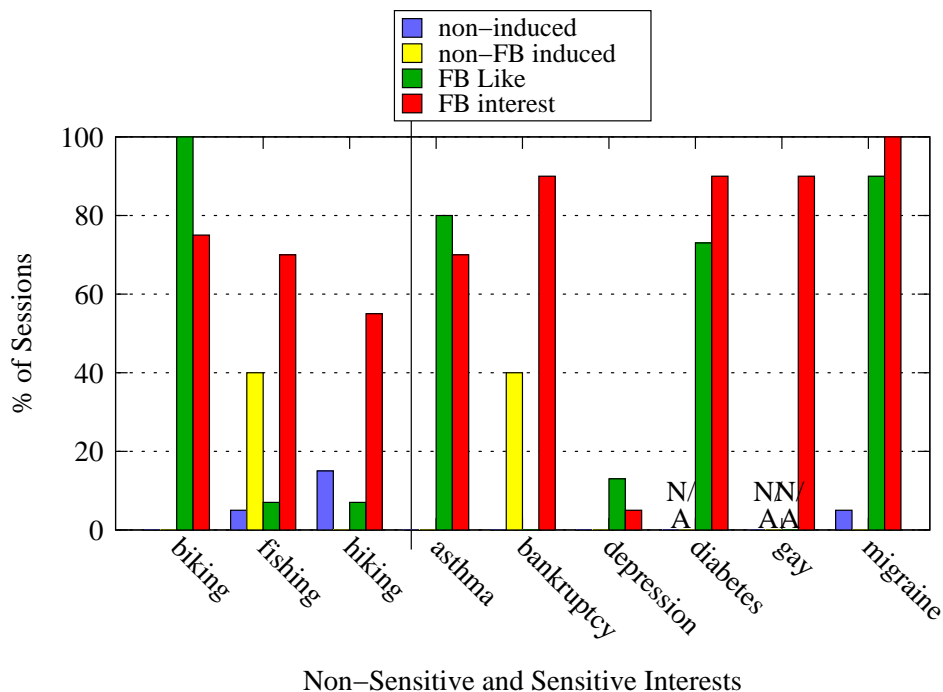


Figure 7: Percentage of Sessions Displaying Facebook Ad Matching Induced Interest

7.2 Facebook Ads for Sensitive Topics

We also looked more closely at how Facebook handles ads for sensitive topics. This question is relevant as the Facebook policy for ads [7] says “Ad text may not assert or imply, directly or indirectly, within the ad content or by targeting, a user’s personal characteristics within the following categories: race or ethnic origin; religion or philosophical belief; age; sexual orientation or sexual life; gender identity; disability or medical condition (including physical or mental health); financial status or information ...”

The results in Figure 7 show a high percentage of sessions where ads matching sensitive interests are shown including ads for diabetes testing (Figure 20) and for a gay friendly real estate firm (Figure 21). Despite targeting sensitive topics, these examples are apparently acceptable according to guidelines provided by Facebook for creating ads [8] because they

make statements describing the product or service and not necessarily the characteristics of the user.

However other ads that we observed do appear to violate Facebook’s own guidelines. One ad showed the user’s age of 32 (Figure 22). Another encouraged recipients to join gay men in creating their own roommate listing (Figure 23). Yet another asks recipients “do you have diabetes?” (Figure 24). These ads assert sensitive information directly or indirectly through language such as asking if a user has a sensitive condition or encouraging them to join others with a sensitive characteristic implying that the recipient has this characteristic. We found ads asserting a sensitive characteristic in *each* of the experimental sessions where the interest was induced as a Facebook interest, primarily through ads for diabetes, migraines and sexual orientation.

8 Summary and Future Work

In summary, our initial study of a few ad networks and our focused study of the Google ad network found many expected contextual, behavioral and location-based ads along with combinations of these types of ads. We also observed some profile-based ads. We generally found that behavioral ads based upon induced interests were shown as categories in the Ad Preference Manager of the ad network, but found a couple unexpected cases where the interests were not visible in the APM. We also found unexpected behavior for the Google ad network in that non-contextual ads were shown related to induced sensitive topics regarding sexual orientation, health and financial matters. However, we also found such ads displayed when the sensitive topics were not induced meaning that Google may not be showing behavioral ads for these topics, but users with such sensitive topics may be unable to discern the difference.

In a smaller study of Facebook, we did not find clear evidence that a user’s browsing behavior on non-Facebook sites influences the ads shown to the user on Facebook, but we did observe such influence when the Facebook Like button is used to express interest in content. We did observe Facebook ads appearing to target users for sensitive interests with some ads even asserting such sensitive information, which appears to be a violation of Facebook’s stated policy.

Our work has a number of directions for future work. Automating the data collection will allow for more experiments and allow them to be done in parallel to reduce the impact of ad churn. We plan a longitudinal study as the results reported here may change over time. We also plan to evaluate the effectiveness of various privacy protection measures.

Addendum

We disclosed results of our study to Google and Facebook as possibly being inconsistent with their stated policies. We did not receive any response to our disclosure.

References

- [1] Alexa: Most popular web sites. <http://www.alexa.com/>.
- [2] Julia Angwin. The web's new gold mine: Your secrets. *Wall Street Journal*, July 30 2010. <http://online.wsj.com/article/SB10001424052748703940904575395073512989404.html>.
- [3] AOL AdVisibility - my advertising preferences. <http://www.advertising.aol.com/advisability>.
- [4] Rebecca Balebako, Pedro Leon, Richard Shay, Blase Ur, and Lorrie Faith Cranor. Measuring the effectiveness of privacy tools for limiting behavioral advertising. In *Web 2.0 Workshop on Security and Privacy*, May 2012.
- [5] Bluekai registry. <http://www.bluekai.com/registry> .
- [6] Compete. <http://www.compete.com/>.
- [7] Facebook advertising guidelines, Last Revised: March 20, 2012. <http://www.facebook.com/adguidelines.php>.
- [8] Facebook help center: Ad and sponsored stories copy, image, targeting, and destination. <http://www.facebook.com/help/?page=245316378826196>.
- [9] Fiddler web debugging proxy. <http://www.fiddler2.com/fiddler2/>.
- [10] Google Ads Preferences. <http://www.google.com/ads/preferences/>.
- [11] Interest-based advertising: How it works - Google ads preferences, 2010. <http://www.google.com/ads/preferences/html/intl/en/about.html>.
- [12] Google privacy policy, Last Modified: March 1, 2012. <http://www.google.com/intl/en/policies/privacy/>.
- [13] Saikat Guha, Bin Cheng, and Paul Francis. Challenges in measuring online advertising systems. In *Proceedings of IMC*, November 2010.
- [14] Balachander Krishnamurthy, Konstantin Naryshkin, and Craig E. Wills. Privacy leakage vs. protection measures: The growing disconnect. In *Proceedings of the Web 2.0 Security and Privacy Workshop*, pages 1–10, Oakland, CA USA, May 2011.
- [15] Balachander Krishnamurthy and Craig E. Wills. On the leakage of personally identifiable information via online social networks. In *WOSN*, 2009.
- [16] Balachander Krishnamurthy and Craig E. Wills. Privacy diffusion on the web: A longitudinal perspective. In *Proceedings of the World Wide Web Conference*, 2009.
- [17] Jonathan Mayer. Tracking the trackers: Where everybody knows your username, October 2011. <http://cyberlaw.stanford.edu/node/6740>.
- [18] Jonathan R. Mayer and John C. Mitchell. Third-party web tracking: Policy and technology. In *Proceedings of IEEE Symposium on Security and Privacy*, May 2012.

- [19] Quantcast. <http://www.quantcast.com/>.
- [20] Franziska Roesner, Tadayoshi Kohno, and David Wetherall. Detecting and defending against third-party tracking on the web. In *Symposium on Networked Systems Design and Implementation*, April 2012.
- [21] Arnold Roosendaal. Facebook tracks and traces everyone: Like this! Technical report, Tilburg Law School, November 2010. Legal Studies Research Paper Series No. 03/2011. <http://ssrn.com/abstract=1717563> .
- [22] Yahoo! ad interest manager. http://info.yahoo.com/privacy/us/yahoo/opt_out/targeting/details.html.

Appendix

This appendix shows a gallery of illustrative examples captured during testing and referenced in the paper.

```
GET http://ad.doubleclick.net/pfadx/pand.default/...;  
  artist=S1421673;genre=love;ag=32;gnd=1;zip=11201  
Host: ad.doubleclick.net  
Referer: http://www.pandora.com/  
Cookie: id=223d4200013312||t=1292486411|et=730|cs=p999khn4
```

Figure 8: Leakage of Private Information from Pandora Profile to DoubleClick



Figure 9: Contextual Ad on nytimes.com Science Page



Figure 10: Location-Based Ad on pandora.com



Figure 11: Contextual and Location-Based Ad on gaylife.about.com



Figure 12: Profile-Based Ads on pandora.com using Profile Location (New York) as well as Profile Age (32) and Current Location (Worcester)



Figure 13: Behavioral Ads for Interests Golf and Dogs (Each Shown in Respective APM) on accuweather.com and macmillandictionary.com



Figure 14: Behavioral Ad for Miami Interest (Not Shown in APM) on pandora.com



Figure 15: Non-Contextual Ad for Sensitive Induced Interest Depression on nytimes.com



Figure 16: Non-Contextual Ad for Sensitive Induced Interest Bankruptcy on accuweather.com Page for Miami Weather



Figure 17: Non-Contextual Ad for Sensitive Induced Interest Pregnancy on pandora.com



Figure 18: Non-Contextual Ads for Sensitive and Non-Sensitive Interests Matching YouTube-induced Interests on thefreedictionary.com



Figure 19: Facebook Ad Based on a User's Interest



Figure 20: Facebook Ad Targeting Health Condition of Diabetes



Figure 21: Facebook Ad Targeting Sexual Orientation



Figure 22: Facebook Ad Asserting Age



Figure 23: Facebook Ad Asserting Sexual Orientation



Figure 24: Facebook Ad Asserting Health Condition of Diabetes