

AdReveal: Improving Transparency Into Online Targeted Advertising

Bin Liu*, Anmol Sheth[‡], Udi Weinsberg[‡], Jaideep Chandrashekar[‡], Ramesh Govindan*

[‡]Technicolor *University of Southern California

ABSTRACT

To address the pressing need to provide transparency into the online targeted advertising ecosystem, we present *AdReveal*, a practical measurement and analysis framework, that provides a first look at the prevalence of different ad targeting mechanisms. We design and implement a browser based tool that provides detailed measurements of online display ads, and develop analysis techniques to characterize the contextual, behavioral and re-marketing based targeting mechanisms used by advertisers. Our analysis is based on a large dataset consisting of measurements from 103K webpages and 139K display ads. Our results show that advertisers frequently target users based on their online interests; almost half of the ad categories employ behavioral targeting. Ads related to *Insurance*, *Real Estate* and *Travel and Tourism* make extensive use of behavioral targeting. Furthermore, up to 65% of ad categories received by users are behaviorally targeted. Finally, our analysis of re-marketing shows that it is adopted by a wide range of websites and the most commonly targeted re-marketing based ads are from the *Travel and Tourism* and *Shopping* categories.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Measurement Techniques

General Terms

Design, Measurement

1. INTRODUCTION

In recent years, online advertisers have attempted to improve the relevancy of ads shown to users by profiling users' online interests and delivering ads relevant to those interests. A recent study [16] showed that online trackers are ubiqui-

tous and cover a large fraction of a user's browsing behavior, enabling them to build comprehensive profiles of their online interests. This widespread tracking of users and the subsequent personalization of ads have received a great deal of negative press; users associate adjectives such as *creepy* and *scary* with the practice [18], primarily because they lack insight into how their data is being collected and used.

Our paper seeks to provide *transparency* into the targeted advertising ecosystem, a capability that has not been explored so far. We seek to enable end-users to reason about why ads of a certain category are being displayed to them. Consider a user that repeatedly receives ads about cures for a particularly private ailment. The user currently lacks a way to reason about which one of the following three targeting mechanisms caused the ad to target her. Is it because the user's online interest profile matches the profile of users the advertiser is seeking to target (behavioral targeting)? Is it because the websites that the user visits are contextually relevant to the ad and draw users that the advertiser is interested in targeting (contextual targeting)? Or is it because the user actually tried to buy the particular medication online previously and the advertiser is re-marketing the product (re-marketing based targeting)? Our position is that providing transparency into which one of the above targeting mechanisms was used would lead to a new class of ad control mechanisms that enable end-users to exert fine-grained control over targeted advertising. Specifically, end-users would be able to block tracking along the actions related to *specific* ads, or indicate their ad preferences at a granularity that is not feasible via existing tools such as Adblock [1] and NoScript [14].

The primary challenge in providing transparency is to design mechanisms that account for the inherent complexity involved in ad delivery. Advertisers can select one or more of the three targeting mechanisms described above, and multiple ad campaigns can co-exist. Consequently, at any point, a webpage could contain ads from multiple campaigns that are targeting different aspects of the user's online interests. Furthermore, the ad selection process is based on a real-time auction, whose outcome also depends on financial parameters of the ad campaign like the cost per mille/thousand impressions (CPM) and desired click through rate (CTR).

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

Hotnets '13, November 21–22, 2013, College Park, MD, USA.

Copyright 2013 ACM 978-1-4503-2596-7 ...\$10.00.

To address these challenges, we have developed *AdReveal*, a practical measurement and analysis framework that relies on end-user measurements to provide transparency into how online display ads (flash and image based ads) are targeted to the user. *AdReveal*'s measurement tool is a first of its kind browser-based extension that provides detailed measurements of online display ads. *AdReveal*'s analysis component uses a novel *contextual model* to predict the ad categories expected on a webpage (in the absence of tracking) and a metric to quantify the extent to which the user is being behaviorally targeted. The contextual model is learnt from a large measurement dataset, obtained using *AdReveal*'s measurement component, which consists of a total of 103K webpages and 139K display ads by simulating the web browsing behavior of 80 users derived from real user AOL search logs. Finally, for re-marketing based targeting, *AdReveal* provides the user the exact actions in her clickstream that led to the ad being targeted.

Our evaluation of the transparency mechanisms provides interesting insights into the ad targeting ecosystem. We find that across all the ad categories in our dataset, almost half of the ad categories employ behavioral targeting. Ad categories of *Politics and Government* is the most contextual while ads from the categories of *Insurance, Real Estate and Travel and Tourism* are heavily targeted towards the user's online interests. We find that between 12-65% of the ad categories that a user receives are behaviorally targeted, and there was not a single user in our dataset that received only contextual ads. Our analysis of re-marketing ads shows that ad categories of *Travel and Tourism* and *Shopping* contained the largest number of re-marketing ads.

2. BACKGROUND

To simplify the management of online ad campaigns, several ad campaign management platforms (i.e., ad targeting platforms) have emerged (e.g., Google AdSense [2]). These enable advertisers to configure ad campaigns by describing the user demographic and interests they wish to target, along with preferences on webpage categories where their ads should appear. These "filters" are selected from a category taxonomy (each vendor defines their own). Users' online interest profiles are inferred by analyzing the webpage history; this is done by tracking users' across different webpages and domains with tracking cookies belonging to the ad network. Typically, a user interest profile is represented as a set of interest categories, and some vendors structure these as a hierarchy (e.g., *Movies* → *Action Films* → *Superhero films*).

There are three primary targeting mechanisms available to advertisers when setting up an ad campaign, namely *contextual*, *re-marketing* and *behavioral* targeting, and these vary in the level of user information used while selecting an ad.

Contextual Targeting involves matching the ad with the context of the page that it is displayed on (and ignores the visitors interest profile). The targeting is implicit: a car insurance company will place ads on auto-related sites because

it is assumed that visitors to the site are likely to own a car (or want to) and will need insurance. With contextual targeting, we expect visitors with different profiles would broadly see the same kind of ads, and these will match the topic/context of the particular website (or a related category).

Re-marketing is a very specific mechanism used by advertisers to target users who, in the past, have indicated a very specific interest in a particular product (e.g., visiting the product website and shops for said product). Say a user visits a car insurance site, clicks on a link to get a quote, but leaves without finalizing it. The insurance company (via the ad-network) can then place *re-marketing* ads – e.g., insurance discounts – into other websites the user visits (which may be unrelated to cars or insurance) to lure the user back to finish the purchase. Here, the advertiser exploits a very narrow and explicit signal from the user to target ads.

Behavioral Targeting is used to select a few "related" ads from a very large ad catalog that is shown to the user, and this filtering is done based on the user's interest profile, computed by the ad-network (by tracking the user over a long period of time). This mechanism goes beyond the "single domain" aspect of re-marketing, and selects ads that *might* relate to the user's long term online interests. With behavioral targeting, a user might see car insurance related ads on a site about *food & nutrition* simply because the user visited multiple different car insurance related websites. This form of targeting is controversial to some, given that it relies on a detailed analysis of the user's online behavior, and results in ads that may be dissonant with the page being viewed

3. ADREVEAL MEASUREMENT TOOL

Collecting measurements about display ads requires the ability to disassemble the elements of a webpage, identify ad elements and associate these with particular categories. Existing ad monitoring and blacklisting tools – AdBlock [1], Ghostery [5], etc. – work by matching URL patterns embedded in a webpage against a set of blacklist patterns, and cannot look deeper into the element and reason about it. The task is made even more difficult by complex DOM structures, deep nesting of elements, and dynamic JavaScript execution, that is found on a large fraction of pages on the Internet today. We developed a browser-based plugin that addresses these challenges that can reliably extract the ad elements of a page, identify the actual landing pages for the ad-elements, and associate pages and the embedded ads with specific semantic categories.

For each page processed by the tool, it extracts the following: (i) the page URL and semantic category, (ii) the destination landing page and semantic category for each display ad, and (iii) embedded re-marketing tags. The tool is implemented as a Chrome browser extension, and the current version is restricted to processing display ads from DoubleClick, which currently has the largest market share for

display ads. In the following, we present a brief overview of the modules of the measurement tool.

DOM Parser/Preprocessor. This module parses the DOM structure of the page and extracts specific attributes of display ads that reveal the landing page for the ad (the website that would be visited by clicking on the ad). This is complicated by the fact that display ads are often embedded in nested `iFrame` tags spanning multiple levels¹. Furthermore, the *same origin policy* enforced by modern web browsers permits an outer `iFrame` to inspect and communicate with its immediate inner `iFrame` only if the two `iFrames` are from the same domain. To address this, we recursively inject custom JavaScript code into all `iFrames` on the webpage and setup a dedicated background page as a communication bridge between nested `iFrames`; this code reads the `<href>` (or `<flashvars>`) attributes for image (or flash) ads, and aggregates this information at the background page running within the context of the plugin.

This module also logs `DoubleClick` elements (re-marketing scripts and cookies) on the page. Re-marketing scripts are detected by searching for the unique `DoubleClick` JavaScript code. `DoubleClick` cookies are detected by monitoring outgoing HTTP requests and comparing against the publicly available patterns provided by the Ghostery [5] tracker database.

Ad Landing Page Extractor. For each identified ad element, this infers the landing page by parsing the value of the attributes extracted by the DOM parser module and searching for specific patterns in the URL like `adurl=`, `redirecturl=`, etc. We manually generate these patterns for `DoubleClick` by inspecting the attribute value. In our experiments, we found that almost 80% of the ads have a landing page that is encoded in these attributes, while the remaining ads require actively following HTTP redirects. We do not follow these redirects; doing so could artificially inflate the click through rates of the ad campaigns, and bias the user profile inferred by `DoubleClick` towards these ad categories.

Semantic Categories of Webpage and Ad URL. *AdReveal* tags every URL (webpage and ad landing pages) with semantic categories, in order to learn the association between them. Since `DoubleClick` does not provide an open API to associate semantic categories to URLs, we use the Yahoo! Content Analysis API [19], which provides categories based on the HTML metadata tags. For example, the URLs `www.nfl.com` \rightarrow *Sports, American Football*, and `www.webmd.com/cancer` \rightarrow *Health, Medical Conditions, Cancer*. Note that a URL can be associated with more than one category.

4. MEASUREMENT DATASET

To study the efficacy of *AdReveal*, we generated a dataset that consists of over 103K webpages and 139K display ads, by simulating the behavior of 80 users from the AOL query log dataset [15]. Note that identifying all the display ad ele-

¹In our experiments we observed up to six levels of nesting.

	Min	Max	Med.	Avg.
Hours to complete the crawl	17.13	79.5	37.75	43.21
Number of pages per user	514	2385	1132.5	1296.4
Number of ads per user	383	2530	944.5	1014.9
Number of ads per page	0	15	0	0.7829
Number of page categories per user	77	164	126	123.6
Number of ad categories per user	60	125	91	90.4375

Table 1: Summary of the dataset used to evaluate *AdReveal*’s transparency mechanisms.

ments in a page requires the page to be downloaded and fully rendered inside a browser context.

The AOL dataset consists of a set of search queries along with a user identifier, and we used this to generate our dataset as follows: we submit each user’s search query to the Bing search engine, and then “visit” the top 5 results returned (with a “reading time” delay of 2 minutes between consecutive visits) – this was done to simulate a user looking for information on a particular topic on Bing. Importantly, we repeat the process twice (with and without history). In the first run, we start with an empty user profile (all cookies cleared), and visit the pages using the default browser settings (allow cookies, cache pages, track history); this “user” is likely to see ads, starting at some point, that target the interest profile (over the pages visited previously). In the second run, we ensure that the profile is blank (no cookies, history or caching), before visiting each page. We refer to these two datasets as the *tracking* dataset and *no-tracking* dataset (respectively).

The *no-tracking* dataset offers a view of what kinds of ads would have been selected if the ad-network had no information about the user. This allows us to build a model that predicts the ad categories that may appear on a webpage. This model is then applied on users in the tracking dataset that enables us to reason about how a user is being tracked (by comparing predictions against the ads being loaded).

Table 2 summarizes the dataset. On average, data for each user took 1.8 days to crawl, yielding roughly 1296 pages with an average of 1015 `DoubleClick` ads within these pages. Each page has an average of 0.78 `DoubleClick` ads, and about 40% of the pages contain at least one `DoubleClick` ad. Out of the almost 104K webpages we browsed, 56% are tracked by `DoubleClick`. The *tracking* dataset contains 81K ads pointing to 9763 distinct landing pages, and the *non-tracking* dataset contains 58K ads, directing to 9073 distinct landing pages. We also found that approximately 9% of the URLs browsed by each user contained re-marketing scripts from `DoubleClick`.

In our dataset, there are 300 distinct semantic categories associated with the web pages and ad landing pages in total. The simulated AOL users in our dataset have a rich and diverse browsing profile, with some categories being more dominant than others. The average number of categories used to describe a user’s browsing profile is 124 while the number of categories to describe the ads is 81. This indicates there is no one-to-one correspondence between webpage and ad categories, and a single ad category may appear across different webpage categories.

Avoiding Measurement Bias. The two datasets are crawled

from the same IP address and within the same time period, eliminating the influence of location based ad targeting and temporal ad campaign settings. We verified that the experiment duration is sufficient for Google to infer a user’s interest profile. We crawl the user profile generated by Google [6] every 10 webpages and observed that user’s interests are updated frequently – for more than 80% of the users the first profiling update appears within the first 10 webpages (with max/avg/median of 30/12/10) and interest categories are added or removed as the profile evolves. Unlike previous work [11], it is sufficient for *AdReveal* to visit a webpage once. This is because we focus on the semantic categories of ads and successive visits to the same webpage with an empty user profile leads to ads served from the same semantic categories. Finally, we validate that the semantic categories used by *AdReveal* and DoubleClick are similar. Across all the users, on average 76% of the semantic categories that describe a user’s browsing history using the Yahoo! API are observed in the user profile generated by Google [6]. The categories not observed in the Google profile are mostly health-related, which Google explicitly avoids [3]. This strong similarity ensures that our analysis is not biased by the Yahoo! semantic categories.

5. ADREVEAL TARGETING INFERENCE

To characterize and measure the different kinds of targeting we need to infer whether (and to what extent), a given ad on a web page is the result of one of the mechanisms discussed in Section 2. In this section, we describe the methodology by which we make that determination based on the annotations provided by the ad measurement tool.

5.1 Analyzing Interest Based Ads

Detecting that an individual ad results from contextual or behavioral targeting is difficult – ad targeting is impacted by dynamic auctions and ad campaign constraints. However, we may be able to infer the extent to which each of these targeting mechanisms are used, at a coarse level, over a *set* of ads belonging to the same category. We develop a metric called *targeting score* that determines which one of the two mechanisms is predominant over a set of ads.

To compute this, we develop a set of binary classifiers (one for each ad category) – denoted as the *contextual model* – which relates an ad category to its associated webpage categories. Recall that these models are required because unlike search ad campaigns, display ad campaigns are setup using semantic categories of webpages. Additionally, there is no one-to-one correspondence between webpages and ad categories (Table 2 provides examples). The targeting score indicates the extent to which the contextual model correctly predicts ad categories on webpages the user visits.

Modeling contextual ads. We train models on the contextual *no-tracking* dataset. We filter out all re-marketing ads from this dataset. Specifically, for a given page that matches a set of web page categories, the model outputs a 1 or 0

Ad Category	Associated Webpage Categories	AUC
Parenting	Parenting, Arts and Crafts, Family Health, Holidays and Celebrations, Education, Cultural Groups.	0.83
Celebrities	Skin Care, Arts and Entertainment, Real Estate, Autos, Celebrities	0.77
Health	Health, Nutrition, Insurance, Disease and Medical Conditions, Athletics, Track and field, Parenting	0.75
Financial Fraud Prevention	Employment and Career, Finance, Credit	0.70
Credit	Finance, Arts and Entertainment events, Credit, Travel organizations, Real estate, Shopping	0.59

Table 2: Model trained for ad categories. The second column enumerates the set of most influential page categories, and third column denotes the AUC score for the model.

for each ad-category, which is a prediction on whether an ad of that category should appear on that web page. If the prediction of an ad appearance holds (output = 1), then the ad-categories *conform* to the trained model, and the targeting is contextual. We experimented with several learning methods, testing each with 5-fold validation. We find that models learned by logistic regression with L1 regularization performed best. These learn a set of coefficients which weigh the relevance of each (input) page category to the (output) ad category. The L1 regularization enforces a sparse model, *i.e.*, only a few coefficients will be non-zero, which fits our input data well, as most webpages are mapped to only a few categories. Importantly, the learned classifier also outputs a confidence score (over its classification results), and we utilize this to account for noise inherent in the data being modeled.

To characterize the models generated using the above described approach, we compute the Area Under the Curve (AUC) score² as well as inspect the model parameters. Across all the 81 ad categories³, the median AUC score is 0.71. 10% of the ad categories had an AUC score above 0.85 (*e.g.*, *American Football*, *Travel Transportation* and *Disease & Medical Conditions*) and 9% of the ad categories had an AUC score below 0.6 (*e.g.*, *Credit*, *Gaming* and *Lottery*).

Table 2 enumerates the most influential webpage categories along with the mean AUC score for five different ad categories. We observe that advertisers make use of diverse but related categories to target contextual ads. In addition to targeting webpages that are semantically related to the ad (*e.g.*, Parenting ads on webpages about *Parenting* and *Family Health*), advertisers also target related webpage categories that draw visitors that the advertiser is interested in targeting (*e.g.*, *Holidays and Celebrations*). Interestingly, we find that ad categories that correspond to low AUC scores are either very broad in scope, *e.g.*, *Arts and Entertainment*, or tend to target a broad spectrum of webpage categories, *e.g.*, *Credit*.

Targeting score: Applying the learned (contextual) model to a user’s web trace (*tracking dataset*), we look at two cases

²AUC scores typically range from 0.5 (random) to 1 (perfect precision and recall)

³We train a model for ad categories with a support of 50 ads.

(in each web page instance): (i) the *true-positive* (TP) case, which validates the classifier prediction and indicates that the ad was selected purely based on the page context, and (ii) the *false negative* (FN) case, where the prediction is incorrect, indicating that the ad was selected based on factors beyond the context of the page (i.e., which the model completely accounts for). The other two cases (true negatives and false positives) are not strong indicators of the ad selection being contextual. Putting these together, we denote the false negative rate, computed for a set of pages, ($FNR = \frac{FN}{FN+TP}$) as the *targeting score*. When *FNR* is close to 0, we expect that the ads placed on the page were contextual; values close to 1 indicate that the ads (behaviorally) target the user.

Improving model robustness. The contextual models trained on the no-tracking dataset have some inherent noise: differences across ad campaigns for the same category and inherent dynamic of ad auctions and campaigns. These serve to weaken the association between webpage categories and the predicted contextual ad category. The noise reduces the confidence of the classifier for the two output classes, and this results in the distributions (of confidence) overlapping. In our models, we found that the overlap is present, but relatively small. Thus, we only consider samples whose classification confidence is above a certain threshold. We omit a detailed discussion due to the lack of space.

5.2 Detecting Re-Marketing Ads

Re-marketing ad campaigns require advertisers to tag different pages on their sites with specific Javascript code generated by the ad platform. This allows the advertiser to distinguish users that reach different parts of their site, and customize the advertising strategies accordingly (e.g., display ads for travel tickets to a specific destination based on the fare search the user performed). Hence, re-marketing campaigns ignore the user profile and “follow” the user on the web, re-marketing the product to convince the user to come back to the advertiser’s page.

AdReveal provides a simple approach to detect re-marketing ads. It monitors and logs all domains visited by the user that embed JavaScript re-marketing code in the page source (these patterns are publicly available for the AdSense targeting platform). Subsequently, for every ad, the domain of the ad landing page (the site that would have been visited if the ad was clicked) is matched against the set of domains containing the re-marketing scripts. When the two match, *AdReveal* can point the user to exact pages in her clickstream that caused the specific ad to be targeted. This enables users direct feedback about how their particular actions in the past results in the current ad.

6. TARGETING CHARACTERIZATION

6.1 Interest Based Ads

We use the *targeting score* metric to characterize the fol-




Figure 1: Cumulative distributions for the targeting score across all categories and two example categories of *Health* and *Travel & Tourism*

lowing aspects of ad targeting: (i) Are there particular categories that predominantly target users contextually (or behaviorally)? (ii) What fraction of ads (and categories) are behaviorally targeted towards a “general” user? (iii) How does the targeting score (the level of behavioral targeting) evolve over the duration of an individual user’s browsing habits (as an interest profile is constructed by the ad networks)?

Targeting Bias. Figure 1 shows the cumulative distribution of the targeting scores across all users and categories in our dataset (All Categories line). This distribution is computed by taking the median targeting score for each category across all users in our dataset. We observe that 53% of the ad categories have a targeting score above 0.5, indicating that almost half of the ad categories contain more behaviorally targeted ads than contextually targeted ads. The most contextual ad category was *Politics and Government* with a targeting score of 0.17 and categories like *Insurance*, *Real Estate* and *Travel and Tourism* primarily employ behavioral targeting with a score of 0.7 and above. Additionally, the graph also shows the distributions across all users for the categories of *Health* and *Travel & Tourism*. The *Health* ad category is highly contextual, with 78% of the users having a targeting score below 0.5⁴. On the other hand, ads about *Travel & Tourism* are primarily behaviorally targeted with 75% of the users having a score above 0.5.

Number of Behavioral Ad Categories Per User. We now focus on characterizing the fraction of ad categories that are behaviorally targeted in a typical users browsing profile. A targeting score for an ad category above 0.5 implies that the user receives more behavioral ads than contextual ads. Surprisingly, we find that behavioral targeting is common and covers a significant portion of the ad categories that are targeted at the user. Our dataset did not contain a single user that received only contextual ads (targeting score ≈ 0). We find that half of the users have more than 50% of the ad categories being behaviorally targeted, with the maximum and minimum of 65% and 12% respectively. With a more conservative targeting score threshold of 0.7, we find that half of the users have more than 35% of ad categories being behaviorally targeted.

Evolution of targeting score. To characterize how the tar-

⁴This conforms with DoubleClick policies that place strong restrictions over health-based behavioral targeting.




Figure 2: Change of targeting score over the number of visited pages for 2 sample users in the *Shopping* and *Health* categories.




Figure 3: Distribution of different domain categories that re-target users.

getting score for a user evolves over time, Figure 2 shows the score for ads from *Shopping* and *Health* categories for two different users (we select two different behaviors). Since we visit each page in the user’s browsing trace every 2 mins, the x-axis spans approximately 300 minutes in time. Even within this short time span, users have very different targeting scores. During the first 40 pages, U1 receives only contextual ads; however, as she continues to visit pages that are not related to *Shopping*, she continues to receive *Shopping* related ads causing the targeting score to increase substantially to 0.8. On the other hand, U2 is served with mostly contextual ads about *Health*, which aligns with our previous observation that *Health* is an overall contextual ad category.

6.2 Re-Marketing Ads

We now characterize the transparency *AdReveal* provides for re-marketing ads. We focus on the aggregate results across all users since the detection of re-marketing ads is straightforward, and ads that a user receives are from websites she visited that contain the re-marketing scripts.

We begin by first characterizing the different ad categories that employ re-marketing. Across our entire tracking dataset, there are 244 distinct webpage categories that contain the re-marketing JavaScript code (81% of the complete set of categories). We found that almost 90% of the webpages that contain a re-marketing script did not result in ads delivered to the users. This could be because not all re-marketing campaigns are *active* or are triggered only after a specific duration after the user visited the specific webpage.

Figure 3 shows for 15 different categories, the number of webpages that contain a re-marketing script and the number of ad impressions the user received. We observe that re-marketing spans multiple categories, including *Health*, *Music* and *Education*, indicating that re-marketing is adopted by a wide spread of online businesses. Since re-marketing is the

most expensive form of targeted advertising, we expect that businesses that generate high revenue from users will employ re-marketing more than others. Indeed, the *Travel and Tourism* and *Shopping* were among the top ad categories that had the largest number of ad impressions.

To further characterize the extent of re-marketing on a different dataset, we crawl the top 100 websites from Alexa across the categories of *Travel*, *Shopping* and *Health* and count the number of websites that contained a re-marketing script. We found again that re-marketing is quite prevalent and was detected in 31% of the *Travel* websites, 28% of the *Shopping* websites, and 13% of the *Health* websites.

7. RELATED WORK

Existing work that seeks to address some of the transparency properties provided by *AdReveal* fall short in several ways. A common approach pushed forward by the industry is the AdChoices [7] initiative and Google’s ad preferences dashboard [6]. These approaches provide users visibility into their advertising profile and allows one to opt-out of certain “categories” across a few online trackers. However, even with the limited participating entities, the mechanisms are not evenly implemented and often hard to use [13].

Various browser tools, such as Ghostery [5], AdBlock [1], NoScript [14] and Collusion [4] provide users visibility into the presence of third-party trackers on websites. However, these tools cannot reason about specific targeting mechanisms employed and consequently provide a very coarse grained control, by either turning off or on all ads and tracking. Policy proposals like Do Not Track [9] provide a regulatory framework over the tracking and profiling of user data but do not enforce compliance.

Finally, a number of privacy preserving targeted advertising solutions require re-factoring large parts of the ad ecosystem. Privad [12] and Adnostic [17] rely on local caching of ads and generation of the user profile, ObliviAd [8] relies on a new secure processing hardware, and RePriv [10] provides browser specific tools for third-parties to extract user profiles from the browser.

8. FUTURE WORK

As part of future work we plan on extending *AdReveal* to a distributed setting to account for IP geo-location and inferred demographic based targeting. Additionally, we will extend *AdReveal*, to empower end-users with a new class of ad control and selection mechanisms that can selectively block certain ad categories. *AdReveal* can identify the likely webpage categories associated with the ad category that is considered private by the user and re-generate a new cookie profile. This new profile is created by a combination of user profile obfuscation achieved by introducing additional interest categories and/or selectively replaying the user’s web history while skipping over sensitive webpages and blocking tracking on future visits to those webpage categories.

9. REFERENCES

- [1] Adblock Plus. <http://adblockplus.org/>.
- [2] AdSense. www.google.com/adsense/.
- [3] Advertising Policies Help - Healthcare and Medicines. https://support.google.com/adwordspolicy/answer/176031?hl=en&ref_topic=1626336.
- [4] Collusion. <http://www.mozilla.org/en-US/collusion/>.
- [5] Ghostery. <http://www.ghostery.com/>.
- [6] Google Ad Preferences Manager. www.google.com/ads/preferences.
- [7] The Self-Regulatory Program for Online Behavioral Advertising. <http://www.aboutads.info>.
- [8] M. Backes, A. Kate, M. Maffei, and K. Pecina. ObliviAd: Provably Secure and Practical Online Behavioral Advertising. In *IEEE S&P*, 2012.
- [9] Do Not Track: Universal Web Tracking Opt Out. <http://donottrack.us>.
- [10] M. Fredrikson and B. Livshits. RePriv: Re-Envisioning in-Browser Privacy. In *IEEE S&P*, 2011.
- [11] S. Guha, B. Cheng, and P. Francis. Challenges in Measuring Online Advertising Systems. In *ACM IMC*, 2010.
- [12] S. Guha, B. Cheng, and P. Francis. Privad: Practical Privacy in Online Advertising. In *USENIX NSDI*, 2011.
- [13] S. Komanduri, R. Shay, G. Norcie, B. Ur, and L. Cranor. AdChoices? Compliance with Online Behavioral Advertising Notice and Choice Requirements. Technical report, CMU-CyLab-11-005, 2011.
- [14] NoScript: JavaScript/Java/Flash Blocker for a Safe Firefox Experience. <http://noscript.net>.
- [15] G. Pass, A. Chowdhury, and C. Torgeson. A picture of search. In *Proceedings of the 1st international conference on Scalable information systems*, InfoScale '06, New York, NY, USA, 2006. ACM.
- [16] F. Roesner, T. Kohno, and D. Wetherall. Detecting and Defending Against Third-Party Tracking on the Web. In *NSDI*, 2012.
- [17] V. Toubiana, A. Narayanan, D. Boneh, H. Nissenbaum, and S. Barocas. Adnostic: Privacy Preserving Targeted Advertising. In *ISOC NDSS*, 2010.
- [18] B. Ur, P. Leon, L. Cranor, R. Shay, and Y. Wang. Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising. In *SOUP*, 2012.
- [19] Yahoo! Content Analysis API. <http://developer.yahoo.com/contentanalysis/>.