

# Characterising User Targeting For In-App Mobile Ads

Imdad Ullah<sup>\*†</sup>, Roksana Boreli<sup>\*†</sup>, Mohamed Ali Kaafar<sup>\*\*</sup>, and Salil S. Kanhere<sup>\*†</sup>

<sup>†</sup>The University of New South Wales, Australia

<sup>\*</sup>National ICT Australia    <sup>\*\*</sup>INRIA, France

Email: {(imdad.ullah, roksana.boreli, dali.kaafar}@nicta.com.au, salilk@cse.unsw.edu.au}

**Abstract**—Targeted advertising is a growing area of interest in both business and research community. In mobile communications, related research works focus on the collection of user's personal information by the mobile apps, protection against such data collection, and the implications of additional traffic generated by the ads on the mobile device resource utilization. In this work, we present a novel analysis of targeted advertising in the Google AdMob advertising network and show insights about the relevance of Google user profiles, and the categories of apps used, on the in-app ads served on smartphones. We define the classes of ads based on the match between received ads and the app (contextual ads), and Google AdMob user's profile (targeted ads). Our analysis reveals that, for all comparable experiments, the proportion of targeted ads is in all cases higher than the proportion of contextual ads. Moreover, blocking the targeting (disabling targeting in an AdMob user profile settings) results in a significant drop in the number of received ads with some experimental instances receiving no ads at all. Overall, the number of targeted ads is comparatively lower than the number of generic ads. Although this could be partially due to the limited size of ad pools at the time of our experiments, there is also an indication that, although user's information is collected, the subsequent use of such information for ads is still low. We present additional insights on the comparison between Google AdMob and other mobile advertising networks and illustrate the dominance of the former in both the number of ads served and the time during which the ads are displayed to the mobile users.

**Index Terms**—Targeted Ads, Privacy, Mobile Apps, Experiments

## I. INTRODUCTION

Advertising in mobile services and in-app ads present an emerging growth area, in which targeting (or behavioural) advertising is becoming an increasingly important source of revenue for both advertisers and advertising companies. Targeted advertising is based on big data analytics, where user's personal information is collected and processed for the purposes of profiling and targeting.

Despite the industry growth, there is still limited research on characterising in-app ads. Prior research work has demonstrated the large extent to which apps are collecting user's personal information [1], the potential implications of receiving ads to user's privacy [2] and the increased utilisation of mobile device resources [3], [4]. To combat these issues, a number of privacy preserving [5], [6], [7] and resource efficient mobile advertising systems [3], [4] have been proposed. Works on the characterisation of mobile ads have primarily focused on

measuring the efficiency of targeted advertising [8], to examine whether the targeted advertising based on the users' behaviour leads to improvement in the click-through rates. However, so far there have been limited insights about the extent to which targeting is effected in mobile advertising, that will ultimately determine the magnitude of various issues including the loss of privacy.

The measurement of in-app advertising presents a number of challenges. While there are several in-browser measurement methods presented in the literature, such as [5], [6], [9], [10], however, these cannot be directly applied to the evaluation of in-app ads. First, the targeting may be based on a number of factors that go beyond what is used in browser-based ads, including mobile apps installed on the device, the way they are utilised (e.g. heavy gamers may receive specific ads), etc. Second, the classification of ads requires unifying of the mobile market place(s) and traditional online environments, as the ads may relate both to merchant web sites and to other apps that may be purchased and downloaded to the mobile device.

In this work, we present a new methodology for classifying in-app ad traffic and carry out a series of experiments to characterise in-app ads. Our main contributions in this paper include:

- Design of a measurement methodology for collecting and categorising in-app advertisements, including the classification mechanism that unifies web based and app ad categories.
- Characterising the ads collected in our experiments (2500 measurement hours running 100 apps resulting in close to 80,000 ad impressions), including the volume and frequency of ads from various ad networks, and ad display/impression times. We further provide insights about the received ad categories and the prevalence of targeting and context in selected categories. We show that targeted ads (in line with user profiles derived by the analytics service and based on user's data) are more prevalent in all Google Play apps experiments, compared to contextual ads (in line with the corresponding mobile app categories). Despite the prevalence of data collection and profiling, the majority of in-app ads are still generic, i.e. not tailored to either user profile or context. Although this could be partially due to the limited sizes of the ad pools for the profiles and context selected in our

experiments, we argue that this result could also indicate that there is significant potential for further growth of targeting for mobile ads. This is in line with a recent study<sup>1</sup>, showing the increased relevance and untapped potential of mobile advertising.

We note that the low level of targeting results in a (currently) low privacy risk from targeted ads, as user's data, although collected, is not effectively utilised. Regardless, this does not reduce the motivation for further development of privacy preserving ad systems, as the projected growth of the mobile ad industry and increased privacy awareness will necessitate such solutions in the immediate future.

This paper is organised as follows. In Section II, we present background related to the ad network eco-system. In Section III, we describe our real-time experimental setup and the procedure for categorising received ads. In Section IV, we discuss experimental results calculated from the proposed methodology. Finally, we conclude in Section V.

## II. BACKGROUND

The ad network eco-system involves different entities which comprise the advertisers, ad agencies and brokers, ad networks delivering ads, publishers and the end customer to whom ads are delivered [1]. For the case of large publishers, the ads may be served both by the publisher and the advertiser [4]. Consequently, the eco-system includes a number of interactions between different parties.

We identify the workflow of a mobile app requesting an AdMob ad and the triggered actions resulting from e.g. a user click. Figure 1 illustrates the advertising networks' key players and their interactions with a user's mobile device. The mobile device requests an ad from the ad network (AdMob) using the POST method (Step 1) containing device version, model, and the app running on device etc. The ad is then served and shown as an overlay ad in the mobile app (Step 2). After the user clicks on the ad, the landing page is associated with both the Google analytic services (Step 3) and the related Web server (Step 5), where some static objects are also downloaded from the CDN (Step 4).

Both the in-browser and in-app ad delivery mechanisms can be customized to an individual user's profile. For in-browser ads, user profiling is performed by analytic companies [11] using different information such as browsing history, web searches etc., that is tracked using cookies. However, for the in-app ads delivery, AdMob may use different information to infer users' interests and to personalise delivered ads, in particular, the installed applications together with the device identifier<sup>2</sup>.

## III. METHODOLOGY

We now present the methodology used to categorise collected ads and to conduct the measurement experiments. Although our primary interest is in the Google ad network and




Fig. 1. AdMob Ad Presentation Workflow.

Android based devices, we note that the presented methodology is applicable to other networks that serve customisable ads on any mobile device.

### A. Ads Collection

We evaluate the ad delivery mechanism by inspecting the traces of traffic generated when executing apps. This enable us to parse and classify the traffic while distinguishing between ad network signalling (accounting, control) traffic (Step 1 of Figure 1), actual ad traffic (Step 2 of Figure 1), content delivery networks (CDN) traffic (Step 5 of Figure 1) and analytics (tracking and profiling) traffic (Step 3 of Figure 1). Our main focus is on the characterization of received ads, with the methodology presented in the following sub-section.

### B. Ads Categorisation

We define the following categories of ads. First, *Targeted* ads, served corresponding to user's profile that includes interests, age, gender or other characteristics of a selected user, as derived by the advertising company and shown in Ads Preferences Managers (Ads Settings), which on Android based devices can be accessed using the Google Settings app. Then, *Contextual* ads, served according to the applications that the ad is displayed in. We note that prior works have used the definitions of both targeting and context in a broader sense for web based ads [12][9]. However as our primary interest is in evaluating the influence of user profiles and apps, for the sake of simplicity we do not study other factors e.g. location. We further define as *Random* the ads that are served to all smartphones (in our experimental setup) in a specified time period; finally, the ads that cannot be classified in any of these ways are *Generic*. Certainly other Ads categorisation approaches could be chosen. For instance, one can deem ads uniquely served to some specific profiles, as targeted ads, and consider others as generic. We could also vary the location of the mobile phones and consider Location-based ads as additional category. However we note that in this paper, the results calculated in subsequent sections are based on the categories criteria/definitions described in this section.

<sup>1</sup><http://corp.aol.com/2013/08/14/new-research-mobile-advertising-driving-purchases-at-unprecedented/>

<sup>2</sup><http://www.google.com/ads/preferences/html/mobile-about.html>

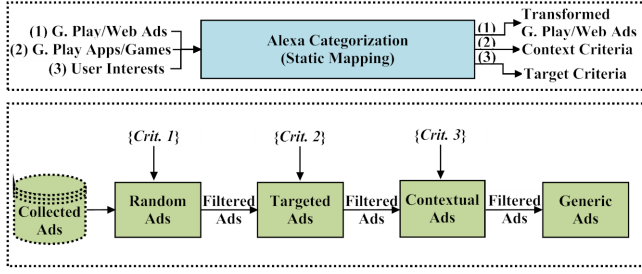


Fig. 2. Ads Categorization Filtering Rule.

We choose the order of filtering for specific classes as shown in (the bottom part of) Figure 2: Random ads are classified and removed from the pool first; then, Targeted ads are classified according to the matching criteria, by which an ad defined by the landing page URL has a high similarity with one or more interest categories in the user profile displayed in Ads Settings. As the next step, Contextual ads are classified according to (again) a match between the app category and the ad (URL). The following sub-section describes the classification methodology used in this paper.

### C. Classification Methodology

The categorisation of ads necessitates matching of the application categories, user profile attributes and ads. An app market place, i.e. Google Play Store, includes a list of all app categories; similarly, the advertising companies publish a list of interest categories that may be included in Ads Settings; finally, ads include the URLs of sponsoring merchants. In order to match those items, we choose a methodology that comprises converting all categories to Alexa web site classification categories<sup>3</sup>. Alexa has 17 major categories which are further divided into a total of 27K sub-categories, to include around 4 Million pre-categorised URLs.

For category matching to Alexa, we represent both the Google and Alexa categories as text strings, and use a standard string matching technique, i.e. we calculate cosine similarity based on tf-idf [13]. Let  $X$  be an Alexa category and  $i$  is a keyword in  $X$ , the  $TF \cdot IDF$  is given as,  $TF \cdot IDF(X, i) = TF_{X,i} \cdot IDF_i$ , where the  $TF_{X,i}$  is the term frequency (occurrence of a particular term) of  $i$  in Alexa category  $X$  and is defined as  $TF_{X,i} = \frac{|t \in X: t=i|}{|X|}$  and  $IDF_i$  is the inverse document frequency of  $i$  in the considered category. In this case the  $IDF_i$  is the logarithm of Alexa category's keywords frequency inverse i.e.  $IDF_i = \log \frac{1}{f_i}$ . These  $TF \cdot IDF$  are computed for every Alexa categories and is saved in a vector. Likewise, the  $TF \cdot IDF$  is computed for user profile, represented as,  $TF \cdot IDF(Y, i) = TF_{Y,i} \cdot IDF_i$ , where the result is stored in another vector. Thus the cosine similarity-based IDF between these two vectors is computed

as

$$\text{Cosine-IDF}(X, Y) = \frac{\sum_{x \in X \cap Y} IDF_x^2}{\sqrt{\sum_{x \in X} IDF_x^2} \sqrt{\sum_{x \in Y} IDF_y^2}}$$

For a resulting cosine similarity of 1, we choose a single Alexa category as a match, while for a lower resulting value we choose a number of matching Alexa categories. Classification of URLs is either obtained directly from Alexa website, or when the URL is Google Play store (for ads related to apps or games); Google Play Store app / game category is mapped to Alexa using the above method. Note that the proposed classification methodology builds on the method proposed in [2].

We also perform a sanity check by manually comparing selected categories, observing that (not surprisingly) there is a high similarity between the top level Alexa categories and Google Play store categories.

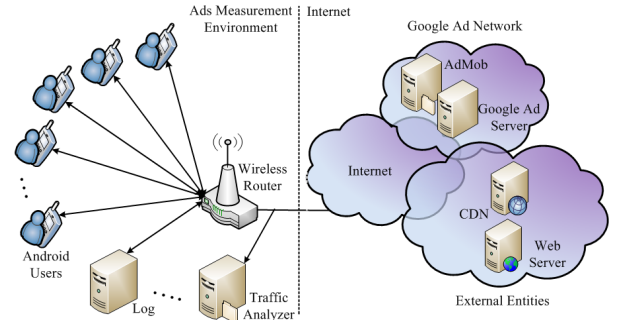


Fig. 3. Network Scenario.

### D. Experimental Setup

We evaluate the in-app Google AdMob ads serving mechanism using the experimental setup shown in Figure 3, with phones having Wi-Fi based network access. Each Android phone is configured (trained) with a selected user profile, corresponding to one of the GooglePlay app categories<sup>4</sup>. Profile training is done by installing the top 100 apps from the selected app category and running the top 10 apps that receive ads (note not all apps receive ads) for a minimum period of 24 hours. During the training process, we note that the phones evolve a sophisticated set of user interests comprising a specific user profile, as shown in the Ads Settings accessed via the Google settings app. A single phone is additionally configured with a Null profile, where the opt-out option (note this applies to profiling and this phone still receives ads) is chosen in the Ads Settings. We stress that our goal is to evaluate the influence of context and user profiles on the received ads.

For ad collection experiments, we install and run the same apps that were used for training, from a selected category (depending on the experiment, the selected category may

<sup>3</sup><http://www.alexa.com/topsites/category>

<sup>4</sup><https://play.google.com/store/apps/>

be different from the category the test phone's profile was configured with). Each experiment is run for a period of 25 hours, on 10 phones in parallel. Out of the total of 27 app categories (including both applications and games) available on GooglePlay, we choose 9 for our experiments: Business, Communication, Education, Entertainment, Games-Arcade & Action, Health & Fitness, Medical, Shopping and Sports. In total, our experiments result in 2500 hours of measurements using 100 apps. Although our test setup was limited by the availability of phones, we consider that the chosen app categories are a sufficiently large representative of the Android advertising ecosystem and that the results presented here can be generalised to the remaining app (and profile) categories.

#### IV. EXPERIMENTAL RESULTS

In this section, we present results obtained according to the methodology outlined in Section III. We first investigate the ad traffic characterization and then detail the categorization of received ads. Our main objective is to quantify the results of in-app ad delivery mechanism and, based on those findings, compute the categorization of received ads.

*a) Ads Characterisation:* During our experiments we receive a total of 79431 ads, of which 65543 are served by Google AdMob, while the remaining 13888 are collectively served from other ad networks e.g. adfonic, admarvel etc. The higher proportion of received ads with Google AdMob can be explained by the fact that, in our experiments, we use applications from Google Play Store.

We first analyse Google AdMob ads and we discuss other ad networks in the following sub-sections. Figure 4 shows CDF for unique AdMob ads with corresponding frequencies (a more detailed view of 80% of unique ads' CDF is also shown), showing the number of times each unique ad is served. It can be observed that 27.14% of unique ads are served once. A higher number of specific unique ads can also be observed. For example, a Google Play ad for NRL (National Rugby League, a popular sports competition) is served 2046 times. Overall, we found that 93% of unique ads are served 14010 times while 7% ads are served 51533 times, which corresponds to a significant proportion (78.63%) of the total AdMob ads. This strongly suggests that ad networks could benefit from the use of caching mechanisms, such as the AdCache technique proposed in [4], to efficiently use network resources by caching the more frequent ads.

Figure 5 shows Alexa categories of received ads along with the frequency bins which refer to the number of times ads are served within the specified frequency, e.g. 1-100 times. Unique ads within a specific frequency bin are also shown on top of each bar. We can observe that ads within the frequency bin 1-100 have a rich set of categories. Across all the frequency bins, we observe that ads from Arts (23.59% of total AdMob ads to the left of figure, i.e. ads within lower frequency bins) and Sports (20.57% of total AdMob ads, within higher frequency bins) categories are served with higher numbers. Note that, the Sports ads primarily relate to Google Play apps NRL and AFL (Australian Football League).




Fig. 4. Frequency (CDF) of Unique Ads.




Fig. 5. Frequency of Unique Ads Served by Categories.

*b) Ads Display Time:* Various parameters we calculate for ads analysis are shown in Figure 6. Comparing Google AdMob with other ad networks, in Figure 7 we show the dominance of AdMob. Considering for example the Google Play Business app category, we can observe a total of 10719 received ads (as indicated on top of 'Business' bar to the left of Figure 7), among which, 91.66% and 8.34% of ads are respectively served by Google AdMob and other ad networks. These ads are displayed for different amount of times, e.g. in the above example, AdMob ads are displayed for 14.3hrs while other ad networks are collectively displayed for 0.96hrs, averaged over all Android phones. Consequently, total amount of time taken by Business apps is 15.3hrs (61.04% of total experimental time of 25hrs) while for a duration of 9.74hrs, no ads are served (*Idle* time). The {Display, *Idle*} time pairs are shown on top of each category of Google Play apps, on the right hand side of Figure 7, e.g. {15.3, 9.7} hours. Thus we observe that, for the entire set of experiments, ads are served for 53.91% of time, whereas the *Idle* time is 46.09%.






Fig. 6. Parameters used for ads analysis in an experiment: *Idle* time  $T_0$ , is the time duration measured from the beginning of the experiment, where no ads are served; the ads Display time is the time duration during which ads are shown on mobile's screen.  $T_G$  and  $T_{NG}$  are, respectively, the time duration or a burst of consecutive ads from Google or other ad networks; Impression time is the aggregate of these individual bursts.




Fig. 7. Comparison of AdMob & Other Ad Networks.

This analysis indicates that ad networks do not sufficiently utilize the available app usage time to serve ads. This suggests that there is still a potential for the Mobile Ads market to grow. We note that, among 900 experimental instances, 310 instances do not receive any ad. Furthermore, among these 310 instances, 70 instances belong to 7 applications that do not receive any ads over the entire set of user profiles. However, these applications do receive ads, as noted a number of times, e.g. on 12:25PM, 30 May, 2013. Possible reasons could be the influence of user profiles, as the Null profile has the lowest proportion of apps that receive ads (Null phone attracts lowest proportion, 7.57%, of total ads). We also observe that, for instance, the Entertainment profile attracts a lower number of ads when the Google Play Business and Education apps are run. This observation strongly suggests influence of the user profiles/interests, related to the presence, absence, or contents of user's profile. Another reason could be that Australia may have a small pool of ads to serve.

c) *Ads Impression Time*: Figure 8 shows the CDF for ad Impression time of Google AdMob and other ad networks. Note that we consider top 5 ad networks, that serve the highest number of ads. The ad Impression time for AdMob, for some experiments, spans the entire duration of the experiment (as shown in inner graph of Figure 8) where no ads are downloaded from other ad networks. Comparative analysis shows that consecutive ads from Google AdMob are more




Fig. 8. Ad Impression Time for AdMob and Top 5 other Ad Networks (CDF).

frequently served compared to other ad networks, with the respective time periods of 3443 and 2847 on the scale of 1000sec. We further observe that the inter-arrival time for both AdMob and other ad networks for a specific experiment are identical.




Fig. 9. Proportion of Random, Targeted, Contextual, & Generic Ads (%).

d) *Categorisation of Ads*: We now discuss the categorization of received ads and we examine the influence of context and user profiles. Figure 9 shows the overall fractions for categories of received ads, classified as Random, Targeted, Contextual, and Generic, for a specific Google Play apps category, according to methodology discussed in Section III-B. The total numbers of received ads (e.g. 9825 ads for the Business category) along with unique ads (185 ads) are also shown on top of each category's bar. Random ads are calculated with overlap time of 6hrs (we also take another variant with a 1hr overlap, where we observe slightly lower number of ads; those results are not shown due to space limitations. It can be observed from Figure 9 that for all Google Play apps, a




Fig. 10. Proportion of Random, Targeted, Contextual, & Generic Ads (%) under different User Profiles.

higher proportion of Targeted ads (corresponding to user profiles) is always served, when compared to Contextual ads. In addition, the percentage of Targeted ads received in Sports and Education apps is greater than the Generic/Contextual/Random ads; likewise, the proportion of Targeted ads for Entertainment apps is slightly lower than that of Generic ads. An important point to note is that there is a higher overall proportion, 77%, of Generic ads in Google Play apps, indicating that Google AdMob is more likely to serve Generic ads, regardless of both forms (according to their profiles or application contexts) of targeting mobile users. Random ads are also distributed with higher percentages, e.g. in Business apps, 18.57% of the ads are randomly served irrespective of user profiles (including the Null profile). Overall, on all the Google Play apps, 11% of the total number of AdMob ads are Random ads.

Figure 10 illustrates the proportion of categories of received ads for all user profiles when only Google Play Entertainment apps were running on the phones. We can observe that Generic ads are served in higher proportion for the majority of user profiles, while in some profiles such as Education and Sports, the proportion of Targeted ads is higher than the Generic ads. The Null profile (having no profile), naturally attracts a higher percentage of Generic ads. We observe similar results for apps from other categories.

While the high proportion of generic ads indicates a relatively low level of profile-based user targeting, we now highlight how having the profile (targeting) enabled still has a significant impact on the served ads. Considering the difference between the ads received on Null profile phones and on phones that have user profiles enabled, we note that Null profile phones receive, on the average, more than 12 times less ads (92.43% of total ads in all our experiments are received on phones with profiles, while 7.5% are received on Null phones).

## V. CONCLUSION

In this paper, we present a comprehensive analysis of the ads serving mechanism of Google AdMob network. We propose

a measurement methodology for collecting in-application ads, and observe the prevalence of targeting and context based ads in selected app categories. However, overall our results suggest that, while there is a significant difference in the number of ads served to users that have enabled profiling (in Google settings) compared to those with no profiles, and a notable proportion of ads is tailored to specific user profiles and applications' context, the level of targeting in the Mobile Ads market is still quite low. Therefore, there is still room for improvement in personalisation and targeting in the Mobile Ads market and systems that will efficiently utilise collected user's data on which such targeting is based. Additionally, the low volume of targeting also indicates that the timing is right for the development of privacy preserving ad systems, as the increased user privacy awareness and evolving regulatory frameworks in a number of countries will make new solutions a necessity for the growing Mobile Ads industry in the immediate future.

## REFERENCES

- [1] I. Leontiadis, C. Efstratiou, M. Picone, and C. Mascolo, "Don't kill my ads!: balancing privacy in an ad-supported mobile application market," in *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*, p. 2, ACM, 2012.
- [2] C. Castelluccia, M.-A. Kaafar, and M.-D. Tran, "Betrayed by your ads!," in *Privacy Enhancing Technologies*, pp. 1–17, Springer, 2012.
- [3] A. J. Khan, K. Jayarajah, D. Han, A. Misra, R. Balan, and S. Seshan, "Cameo: A middleware for mobile advertisement delivery," in *ACM MobiSys*, 2013.
- [4] N. Vallina-Rodriguez, J. Shah, A. Finamore, Y. Grunenberger, K. Paganiannaki, H. Haddadi, and J. Crowcroft, "Breaking for commercials: characterizing mobile advertising," in *Proceedings of the 2012 ACM conference on Internet measurement conference*, pp. 343–356, ACM, 2012.
- [5] S. Guha, B. Cheng, A. Reznichenko, H. Haddadi, and P. Francis, "Privat: Rearchitecting online advertising for privacy," *Proceedings of Hot Topics in Networking (HotNets)*, 2009.
- [6] V. Toubiana, A. Narayanan, D. Boneh, H. Nissenbaum, and S. Barocas, "Adnostic: Privacy preserving targeted advertising," in *NDSS*, 2010.
- [7] H. Haddadi, P. Hui, and I. Brown, "Mobiad: private and scalable mobile advertising," in *Proceedings of the fifth ACM international workshop on Mobility in the evolving internet architecture*, pp. 33–38, ACM, 2010.
- [8] J. Yan, N. Liu, G. Wang, W. Zhang, Y. Jiang, and Z. Chen, "How much can behavioral targeting help online advertising?," in *Proceedings of the 18th international conference on World wide web*, pp. 261–270, ACM, 2009.
- [9] C. E. Wills and C. Tatar, "Understanding what they do with what they know," in *WPES*, 2012.
- [10] R. Balebako, P. Leon, R. Shay, B. Ur, Y. Wang, and L. Cranor, "Measuring the effectiveness of privacy tools for limiting behavioral advertising," in *Web 2.0 Workshop on Security and Privacy*, 2012.
- [11] T. Trzcinski, "Analyse, target & advertise privacy in mobile ads,"
- [12] A. Farahat and M. C. Bailey, "How effective is targeted advertising?," in *WWW*, pp. 111–120, ACM, 2012.
- [13] M. Cunche, M. A. Kaafar, and R. Boreli, "I know who you will meet this evening! linking wireless devices using wi-fi probe requests," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2012 IEEE International Symposium on a*, pp. 1–9, IEEE, 2012.