Treads: Transparency-Enhancing Ads

Giridhari Venkatadri
Northeastern University

Alan Mislove
Northeastern University

Krishna P. Gummadi
MPI-SWS

ABSTRACT

Online advertising platforms such as those of Facebook and Google collect detailed data about users, which they leverage to allow advertisers to target ads to users based on various pieces of user information. While most advertising platforms have transparency mechanisms in place to reveal this collected information to users, these often present an incomplete view of the information being collected and of how it is used for targeting ads, thus necessitating further transparency.

In this paper, we describe a novel transparency mechanism that can force transparency upon online advertising platforms: transparency-enhancing advertisements (Treads), which we define as targeted advertisements where the advertiser reveals information about their targeting to the end user. We envision that Treads would allow third-party organizations to act as *transparency providers*, by allowing users to opt-in and then targeting them with Treads. Through this process, users will have their platform-collected information revealed to them, but the transparency provider will not learn any more information than they would by running a normal ad. We demonstrate the feasibility of Treads by playing the role of a transparency provider: running Facebook ads targeting one of the authors and revealing partner data that Facebook hides from users but provides to advertisers (e.g., net worth). Overall, we believe that Treads can tilt the balance of power back towards users in terms of transparency of advertising platforms, and open promising new avenues for transparency in online advertising.

1 INTRODUCTION

Today, many popular Internet services rely on online advertising as their primary source of revenue. Unlike traditional "broadcast" advertising (e.g., television and newspaper ads), online advertising can be personalized to each user. This has led to the emergence of online advertising platforms—e.g., Facebook, Google, and Twitter—that enable *targeted advertising*: on these platforms, advertisers can choose to show their ads only to particular sets of users, based on advertiser-selected fine-grained targeting parameters. To enable this, the platforms create detailed profiles of their users based on activity and information from both on and off their platform.

Recent events such as the debate over the Facebook data that Cambridge Analytica [3] collected and used in targeting have underscored the need for advertising platforms to be more transparent about (a) the data they collect about users, and (b) how advertisements are targeted. Unfortunately, while some advertising platforms offer transparency mechanisms, they have been found to reveal only an incomplete view to users. For example, Facebook's advertising platform was recently shown [1] to not reveal any user information that is sourced from third parties (e.g., data brokers), despite this information being available to advertisers for targeting [5]; Facebook was also shown [1] to reveal at most one attribute used for targeting, even when the advertiser is allowed to specify any number. Thus, relying only on advertising platforms themselves to provide full transparency is unlikely to present a complete picture to users.

Instead, in this paper we propose a novel approach that we call *Transparency-Enhancing Advertisements* (Treads). Treads bring transparency to advertising platforms from the outside by exploiting the very features that advertising platforms use to provide advertisers with targeted ads. The idea behind Treads is quite simple: Treads are targeted advertisements where the advertiser reveals information about their targeting to the targeted user (e.g., by providing the targeting information in the ad itself). Thus, Treads increase transparency based on the very nature of platforms' functionality: a user is supposed to see a targeted ad if and only if they satisfy the advertiser's targeting parameters. Thus, if a user sees a Tread, it means that the user satisfies the targeting parameters revealed by the Tread.

Treads by themselves do not impact user privacy any more than other ads. In particular, advertising platforms are designed to not reveal to the advertiser *which particular users* satisfy their targeting parameters. Thus, like with other targeted ads, an advertiser cannot know which particular users saw the advertiser's Treads, thus limiting the information learned by the advertiser.

We envisage that Treads will be run by an entity, such as a non-profit, with the goal of revealing to users what information has been collected about them by various advertising platforms. In essence, this non-profit would act as a *transparency provider* by signing up as an advertiser on one or more advertising platforms. Users can then opt-in to receive ads from this non-profit. The non-profit then runs Treads targeting all the opted-in users that satisfy a particular targeting parameter, with different ads for each possible targeting parameter. Each user sees only those Treads corresponding to the targeting parameters they satisfy, and therefore learns

HotNets-XVII, November 15–16, 2018, Redmond, WA, USA

ACM ISBN 978-1-4503-6120-0/18/11...$15.00
https://doi.org/10.1145/3286062.3286089

what these parameters are from the content of the Treads; the transparency provider on the other hand cannot learn what targeting parameters each particular user satisfies.

As proof of concept, we use the above mechanism to reveal data broker attributes (which are not currently revealed by Facebook as mentioned above) of one of the U.S.-based authors of this paper; we successfully reveal eleven attribtutes including information about their net worth. We further show that this approach is very cheap for the transparency provider, typically costing $0.002 per targeting parameter revealed.

Treads potentially violate the Terms of Service (ToS) of the advertising platform, which sometimes disallow the revealing of targeting information explicitly [8, 18, 33]. Typically, these ToS terms are present to avoid "creepy" ads that surprise users; this would not be the case with Treads, as users would have explicitly opt-ed in to receiving them. Additionally, in Section 4 we show how Treads could be potentially run in compliance with ToS by obfuscating the targeting information, or revealing the targeting information on the external landing page, rather than in the ad itself.

Finally, we discuss how Treads could enable *advertiser-driven* transparency mechanisms that are complementary to the current *ad-platform-driven* ones. We hope this leads to regulators pressuring advertisers to be directly transparent to users. For example, advertisers could reveal what they learn about users when the users click on their ads, or could reveal their intent in targeting the ads. While advertising platforms have so far had sole power in determining how transparency mechanisms would be implemented, we believe that Treads will help to balance the scales back in favor of users and regulators.

2 BACKGROUND

We first provide background on online advertising platforms, on their targeting mechanisms, and on their existing transparency mechanisms.

2.1 Online advertising

A key feature that makes online advertising platforms popular with advertisers is their support for targeted advertising, which can take a variety of forms. Advertising platforms typically allow advertisers to define sets of users based on a variety of criteria called *targeting parameters*; the resulting set of users is called an *audience*. Advertisers can then place ads that will only be shown to users in the audience.

Targeting parameters We briefly describe the most salient targeting methods supported by advertising platforms.

Attribute-based targeting Traditionally, platforms have allowed advertisers to target users by their *attributes*[1], often including demographics (age, gender, race), location, employment information, likes and dislikes, political leanings, and even financial or medical information. While some attributes are computed by the platform, others might be obtained through partnerships with third parties; for example, as of early 2018, Facebook's advertising platform provided 614 attributes computed internally by Facebook, as well as 507 additional attributes sourced from data brokers such as Acxiom and Oracle Data Cloud [1]. [2] Finally, many platforms allow advertisers to construct Boolean expressions for targeting, for example allowing advertisers to target all users who match "Millennials who live in Chicago, are interested in musicals, are currently unemployed, and are not in a relationship."

While advertising platforms generally have a pre-selected list of attributes that advertisers can choose from, some also allow advertisers to target users via a wider set of attributes. For example, Facebook's advertising platform allows advertisers to search by particular keywords [1] and select from a list of targeting attributes that match the searched keywords. Google's advertising platform on the other hand allows advertisers to create audiences, called *custom affinity* or *custom intent* audiences [19], where advertisers can specify a series of phrases or URLs that describe the users they want to target, which are then internally used by Google to create an audience of matching users.

Activity-based targeting More recently, all major advertising platforms (including Facebook, Google, and Twitter) have begun to allow advertisers to create audiences of users based on activities that the users have performed (such as visiting the advertiser's website, or using the advertiser's app). These are commonly implemented using either web tracking pixels or advertising IDs from mobile devices.

PII-based targeting Advertising platforms have recently also begun allowing advertisers to specify *exactly* which users they want to target (rather than just specifying their attributes). Such targeting, which we call *Personally Identifying Information (PII)-based targeting* [36], typically requires advertisers to upload a list of PII (such as phone numbers and email addresses) corresponding to the users they want to target; the platform then internally matches these PII to platform users and provides the advertiser with an audience. PII-based targeting is supported by all major platforms including Facebook [39], Google [2], and Twitter [32].

2.2 Transparency mechanisms

Driven by pressure from regulators, many advertising platforms have introduced transparency mechanisms to help users understand the data that is collected and how it is used. While these mechanisms are mostly new and poorly-understood, those that have been studied have been found to provide an incomplete view to users [1, 4].

One common approach is to reveal to users the information through which advertisers can target them, via an "ad

[1] In practice, attributes are typically binary (e.g., "is single" or "is worth between $1M and $2M"). However, attributes such as location or age can take on a range of values (e.g., advertisers can typically target users in a ZIP code, or within a radius around any latitude and longitude).

[2] Owing to privacy concerns, Facebook has recently removed these targeting attributes sourced from data brokers from its advertising platform [31]. It is unclear, however, whether Facebook continues to internally retain attributes sourced from data brokers.

preferences" page. For example, Facebook [15], Google [17], and Twitter [34] all reveal to a user a list of their attributes that an advertiser can use. Facebook and Twitter additionally also reveal to the user a list of advertisers who are using either activity-based retargeting or PII-based targeting to target them [7, 17, 40]. However, none of the platforms reveal to users *which* of their PII can be used to target them, or *which* of their PII particular advertisers are using to target them (which can include data the user did not themselves provide [35]).

Another mechanism, provided by ad platforms including Facebook, Google, and Twitter is to provide with each ad an explanation, *generated by the platform*, about why the ad was shown to the user. Additionally, primarily driven by pressure from lawmakers and regulators to help understand politically-oriented ads [9, 34], ad platforms have also begun to make advertiser activity more transparent on their platforms.

3 TREADS

Transparency-Enhancing Advertisements (Treads), are targeted advertisements in which the advertiser reveals information about their targeting to the receiving user. This targeting information could be included directly within the content of the ad shown to the user by the ad platform (such as part of the ad's text or image or other multimedia content), or could be in one of the landing pages that the links within the ad point to. Further, as shown in Figure 1, this information could either be explicit (immediately readable by humans), or encoded (and thus obfuscated) via some mapping of targeting information to encodings that is provided to users.

For example, an ad targeting people interested in Salsa dancing could simply contain text like "You are interested in Salsa dancing according to this ad platform", or could transform this text into innocuous looking text that users' browsers (e.g., via an extension) know how to decode. Alternately, this information could be encoded into the ad image or other multimedia content (in the ad or in the landing page) via steganographic techniques, which can be extracted by code.

Finally, the targeting information could either be the actual targeting attributes that the advertiser specified on the ad platform, or could be the *intent* of the advertiser (i.e., who they actually wanted to target). The two might differ, for example, in the case of an advertiser who wants to target "experienced professional Salsa dancers" (intent), but is limited by the options on the advertising platform to target "People aged 30 and above who are interested in Salsa dance" (the actual targeting attributes).

3.1 Revealing user information via Treads

We next describe how Treads could make user information collected by advertising platforms transparent to users.

Setup We envision that an entity such as a non-profit could act as a *transparency provider* that aims to help users understand what information has been collected about them by advertising platforms, without seeking to learn this information itself. To this end, the transparency provider signs up as an advertiser on one or more advertising platforms, enabling it to run ads (Treads in particular) to users on those platforms. Since anyone can be an advertiser on most major advertising platforms (e.g., anyone with a Facebook account can be an advertiser on Facebook), potentially any user or entity could act as a transparency provider.

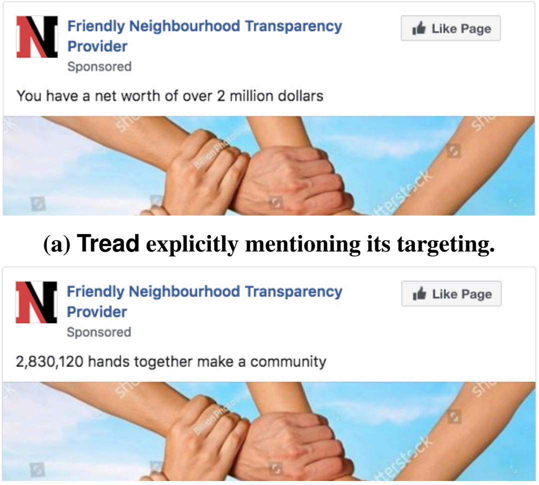

Friendly Neighbourhood Transparency Provider
Sponsored

You have a net worth of over 2 million dollars

(a) **Tread explicitly mentioning its targeting.**

Friendly Neighbourhood Transparency Provider
Sponsored

2,830,120 hands together make a community

(b) **Tread obfuscating its targeting, encoding the parameter as part of the ad ("2,830,120").**

Figure 1: Screenshots showing two different Treads targeting users on Facebook with a net worth of over $2M.

Each Tread the transparency provider runs would reveal one bit of information to the users that it reaches. Thus, for binary attributes like "is single", an ad targeting all users who satisfy that attribute can reveal to the users whether this attribute is set to true. Alternately, a Tread that *excludes* users who satisfy that attribute can reveal to the users that the attribute is either set to false, or is missing from the advertising platform's database for those particular users. For non-binary attributes like location, a Tread can reveal whether the attribute is set to a particular value for the user (e.g., whether a user is determined to have recently visited a particular ZIP code as per the advertising platform). The transparency provider selects a set of such attributes (potentially the preselected set of attributes that the advertising platform offers advertisers), and pays to run one Tread corresponding to each attribute (containing information about that attribute).

Users see these Treads while browsing normally (and can potentially save these using a browser extension). Thus, users learn their information without it being revealed to the transparency provider (by design of advertising platforms); we analyze the privacy of Treads later in the section.

User opt-in While the transparency provider could simply target all users in their country, this might be prohibitively costly and might be undesirable to some users. Thus, instead, the transparency provider could only focus these ads on a set of users who sign-up to the provider's transparency service. If the transparency provider obfuscates Treads as explained above, the provider can share the mapping of targeting information to encodings with users when they opt-in.

Having individual users opt-in is possible because all major platforms now include PII-based targeting or retargeting based on user activities (such as visiting the advertiser's website), as discussed in Section 2. Users could sign-up by providing PII, such as their email address, to the transparency provider (in this case users are not anonymous to the transparency provider). Alternately, in order to remain anonymous to the transparency provider, users could visit a website that the transparency provider owns, where the transparency provider places a tracking pixel provided by the advertising platform.[3] This method has the added benefit that by placing tracking pixels from multiple advertising platforms on the website, the transparency provider could at one shot allow the user to sign-up to learn the information collected about them by multiple advertising platforms. Regardless of the method, the transparency provider can thus run the previously discussed Treads only to an audience of users who opt-ed in.

Validation To demonstrate the utility of Treads, we set ourselves up as a transparency provider, aiming to make Facebook partner categories (attributes on Facebook sourced from external data brokers, but not currently revealed to users) transparent to users. Facebook provides different attributes in different countries; we only focus on those provided to U.S.-based advertisers. We registered as a U.S.-based advertiser on Facebook using a new U.S.-based Facebook account (so as to be unrelated to any other Facebook accounts), and had the two U.S.-based authors of the paper sign-up by liking a Facebook page that we as the transparency provider had created. We ran one ad targeting the signed-up users (i.e., the U.S.-based authors) with each of the 507 binary partner attributes provided by Facebook corresponding to the U.S. We encoded information about the targeting attribute in the text of the ad in an obfuscated manner. Finally, we set the bid cap (the maximum amount that is bid) for each ad to be $10 per thousand impressions or CPM (cost per mille)—five times its default value of $2 CPM for U.S. users— to increase the chances of these ads winning the ad auction conducted by the platform and getting delivered. To test whether the signed-up users were reachable with ads, we ran one control ad where we targeted all the signed-up users without specifying any additional targeting parameters.

While both authors received the control ad, only one author received ads corresponding to his partner categories, receiving eleven different ads corresponding to various attributes including net worth, purchase behavior (particular kinds of restaurants purchased at, particular kinds of apparel purchased), job role, home type, and the kind of automobile they are likely to purchase in the near future. The author who only saw the control ad might not have any information about him collected by data brokers, since he is a graduate student who has only been in the U.S. for over a year. While our validation focused on Facebook, a similar mechanism could be used on other advertising platforms such as Google and Twitter.

Cost The above ads had zero cost since too few users were reached. However, in general the transparency provider must pay the ad platform whenever impressions of Treads are shown to users. The per-user cost of Treads therefore depends on the number of targeting parameters revealed per user. Given the typical recommended bid of $2 CPM from above, each attribute would cost $0.002 to reveal. [4] Thus, using the recommended bid of $2 CPM, it would cost the provider $0.10 to run ads to reveal all targeting parameters to a user who had (say) 50 targeting parameters, showing the financial viability of the mechanism.

Note that there is *zero* per-user cost for running Treads corresponding to targeting parameters that a user *does not* have, as these are never shown to the user. This means that Treads can be cost-effective even for exposing the values of non-binary attributes: for an attribute that can take one of m possible values, the provider would run one Tread targeting each possible value, but would only have to pay for one impression per user, costing around $0.002.

This cost could be paid for by the transparency provider itself (e.g., via donations). Alternately, users opting-in could pay the transparency provider a nominal fee (the cost of their own impressions), making the transparency provider's operations both scalable and sustainable. We leave a full exploration of the funding model to future work.

Scale For a non-binary attribute (such as age) with m possible values, only $log_2(m)$ Treads are required in total to allow any user to learn which of the m possible values they have (since each Tread can represent one of the $log_2(m)$ bits to be learnt). Otherwise, given m binary attributes, m Treads are required to check which ones are set to true for a given user.

Privacy analysis We next describe our threat model and analyze its privacy properties.

Threat model We assume a set of users have opted-in to a transparency provider's service while remaining anonymous to the provider (via the use of a tracking pixel from the advertising platform as previously described). The transparency provider has access to the performance statistics reported by the advertising platform (e.g., for billing purposes); this could include estimates about the number of users reached by different ads. The provider might also be able to associate targeting information with users' cookies (that the provider places on the landing pages); however, we assume there is no further interaction between the user and the provider.

Analysis Given this threat model, the transparency provider can estimate how many of the opt-ed in users have a particular attribute. However, assuming that the advertising platform is designed not to leak the information of individual users

[3] On all major advertising platforms, the identity of users who browse a site with a tracking pixel is not revealed to advertisers; the advertisers are simply allowed to place ads to this group.

[4] For our elevated bid of $10 CPM used in the validation, each attribute would cost $0.01 to reveal.

to advertisers, the transparency provider cannot learn *which* particular users have which attributes.

Besides, assuming that the targeting information is placed within the ad (and not on an external landing page), then the user would not have to leave the confines of the ad platform by clicking the ad, leaving no scope for leakage except via the platform (which we just discussed). In case the targeting information is placed on an external landing page, users can avert any possible leakage by clearing out their cookies and disabling cookies before they start receiving any Treads from the transparency provider.

Revealing a wider variety of information While it is valuable to just make a pre-selected (by the transparency provider) list of attributes transparent as done above, the provider could additionally help reveal a wider variety of information. This could include PII, non-binary attributes such as location, attributes outside the default lists offered by platforms etc. It can be hard for a transparency provider to run Treads corresponding to every possible such targeting parameter that the advertising platform supports; instead, we discuss how the provider could extend the previously described mechanism to accomplish the goal of revealing a wider variety of targeting information.

Supporting PII To enable users to check whether the advertising platform has collected a particular piece of their PII (such as a phone number), the transparency provider could ask users to provide them with PII, and then run a Tread targeting a PII-based audience of all the users who provided them with PII. If a user sees the Tread, it means that the advertising platform has the particular piece of PII they provided the transparency provider. Since advertising platforms generally only require hashed PII to create a PII-based audience [20, 32, 38], the user only needs to provide PII to the transparency provider in hashed form.

Supporting custom attributes To enable users to check whether they have a particular attribute that is not in the pre-selected list of attributes, the transparency provider could have users opt-in to learn such attributes in a custom manner, on a per-attribute level. As with the overall opt-in in Section 3, the users can remain anonymous to the provider when opting-in, as follows: *First*, the transparency provider could have users select an attribute they want to learn, and accordingly redirect them to a distinct (for each attribute) web-page on which they have placed a distinct tracking pixel from the advertising platform. The provider then runs a Tread targeting the audience of visitors to this page (tracked by the ad platform via the tracking pixel, and anonymous to the provider) who also have the corresponding attribute.

4 DISCUSSION

We next discuss challenges and opportunities to the mechanism (Treads) proposed in the previous section.

Advertiser-driven transparency Other than enabling a transparency provider as previously described, Treads also allow any *advertiser* (and not just the advertising platform) to directly include explanations about why they are targeting a particular ad. We hope this leads users and regulators to demand transparency from advertisers as well in the future. For example, advertisers can often learn information about users who click on their ads (e.g., by associating the targeting parameters of the ad with the user's cookie); advertisers could be required to reveal the learnt information to users.

Alternately, advertisers might be required to explain their intent in targeting a particular set of users. Recall the example from Section 3, where an advertiser who wants to target experienced professional Salsa dancers (the intent of the advertiser), is forced to specify due to the limited options on the advertising platform that they want to target "People aged 30 and above who are interested in Salsa". Since any explanation generated by the advertising platform is limited by the targeting parameters specified by the advertiser (which might only approximate the advertiser's intent, as in the above example), an explanation directly revealing the advertiser's intent would complement one generated by the advertising platform. Such an explanation would be even more beneficial in case this advertiser had obtained a list of experienced professional Salsa dancers from some external source (such as a data broker) and used this information to create a PII-based audience; in such a case, an explanation generated by the advertising platform would completely fail to capture the advertiser's intent.

Trusting advertiser-provided explanations It can be hard to directly verify with the advertiser whether their explanations are accurate and complete; however, their explanations might be verified against the corresponding explanations (independently) generated by the advertising platform. Indeed, the existence of these independently generated explanations could force both the platform and the advertiser to provide more accurate and complete explanations.

Co-operation from platforms Treads could potentially violate the terms of service (ToS) of various advertising platforms. Facebook's ToS mention that ads, "must not contain content that asserts or implies personal attributes" [8]. Similarly, Twitter's ToS mention that "Your advertisement must not assert or imply knowledge of personal information" [33], and Google's ToS mention that advertisers cannot run ads that "imply knowledge of personally identifiable or sensitive information within the ad", or that "collect or contain personally identifiable information (PII), unless using an ad format provided by Google and designed for that purpose" [18]. Thus, Treads that explicitly reveal targeting information within the ad itself (and not on a landing page) may violate these ToS.

On the other hand, Treads where the information about targeting parameters is obfuscated would appear to meet the current ToS of platforms, especially if this obfuscated information is placed on an external landing page. In the scenario where regulators require explanations from advertisers via Treads, domains might have no choice but to allow such ads. However, it is beyond the scope of this paper to further analyze under what scenarios Treads pass platforms' ToS.

Legal implications of violating ToS: The legal implications of violating ToS are currently the subject of debate in the U.S. courts [30]. Given that policy is moving towards allowing users the right to know what data has been collected about them by services (e.g., with the enforcement of the GDPR), Treads can be viewed as a user-driven mechanism to perform this task.

Evading shutdown: If advertising platforms forbid all forms of Treads in the future, detection or shutdown of Treads could still be made difficult by distributing them across a number of advertising accounts, effectively crowdsourcing the transparency provider. For example, a number of privacy-conscious organizations or individuals could each create an advertising account and run a few Treads, with each account being responsible for a small subset of the overall set of targeting attributes offered by the platform. This is feasible since anyone can be an advertiser on these advertising platforms.

5 RELATED WORK

We now discuss work that is related to this paper.

Transparency in online advertising platforms There has been a significant amount of effort to improve transparency in online advertising platforms from the outside [23–25, 27, 41]. These approaches work by correlating information about users with the ads that they see, in order to determine whether ads are targeted and how. While these approaches are valuable, they can also be challenging to deploy, requiring either a large diverse population to sign-up (and share their demographic information), or a large number of (fake) control accounts to be created in order to make statistically significant claims. Our approach is complementary to these efforts (and potentially simpler to deploy), and uses the targeting features of the advertising platform *itself* to bring transparency.

Recent work also proposed a mechanism to study which sources of PII are used by Facebook to gather data for its PII-based advertising feature [35], finding that a number of sources including phone numbers provided for security purposes (such as for two-factor authentication), and phone numbers synced from friends' contact lists are used.

Other studies have examined the existing transparency mechanisms offered by advertising platforms. Researchers have studied Google's ad settings page and demonstrated that it does not reveal all the information collected about users [4, 37]; other work has found similar results for Facebook's ad preferences page as well as for the explanations that Facebook provides for why users were shown particular ads [1]. Taken together, these studies motivate the need for additional transparency mechanisms such as ours.

Finally, researchers have audited the accuracy and completeness of the reports provided by Google's advertising platform to *advertisers*, showing that it seems to provide incomplete information to advertisers [26].

Malicious and discriminatory advertising A number of works have found that malicious advertisers could abuse Facebook's advertising platform to infer sensitive information—including PII (e.g., phone numbers) and targeting attributes—about individual users [21, 36]. Targeted advertising platforms have also been subject to other leaks [11, 13, 22]; Facebook has taken steps to address all the above mentioned leaks [16, 28]. While it is hard to say whether advertising platforms might be subject to further leaks of information, for the purposes of this paper, we assume that any such leaks will be patched and that the advertising platform would not leak information about individual users to advertisers.

Facebook's platform has also been shown to be exploitable to launch discriminatory advertising [10, 12], including in covert ways [29]. While Facebook has taken some action to address the problem [14], it was still possible to deploy discriminatory advertisements as of November 2017 [6], which is not surprising given the multiple covert ways of launching discriminatory advertisements that have been found [29].

6 CONCLUSION

In this paper, we proposed a novel transparency mechanism, transparency-enhancing advertisements (Treads), which are advertisements through which the advertiser reveals platform targeting parameters to end users. Given the degree of influence targeted advertising platforms have on people's lives and on countries' politics, and given the disincentives for advertising platforms to be transparent to users, it is essential to break the monopoly that advertising platforms currently have on deciding what information to make transparent.

We showed how Treads can enforce transparency in multiple ways, including by allowing a transparency provider to reveal to users what information about them has been collected by various advertising platforms; we demonstrated in particular how this mechanism could be used to reveal attributes on Facebook that are not currently revealed to users. We also discussed how Treads, by providing a direct channel for advertisers to provide explanations to users, might lead to user demand for such direct explanations from advertisers (about what user information the advertiser learns, or about the intent of the advertiser in targeting particular users). Overall, we believe that Treads have the potential to significantly disrupt the status quo in terms of transparency in advertising platforms, and welcome deployments and further work from the community.

7 ACKNOWLEDGEMENTS

We thank the anonymous reviewers and our shepherd, Ben Zhao, for their helpful comments. This research was supported in part by the Data Transparency Lab and NSF grant CNS-1616234.

REFERENCES

[1] A. Andreou, G. Venkatadri, O. Goga, K. P. Gummadi, P. Loiseau, and A. Mislove. Investigating Ad Transparency Mechanisms in Social Media: A Case Study of Facebook's Explanations. *NDSS*, 2018.

[2] About Customer Match. https://support.google.com/adwords/answer/6379332?hl=en.

[3] C. Cadwalladr and E. Graham-Harrison. Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election.

[4] A. Datta, M. C. Tschantz, and A. Datta. Automated Experiments on Ad Privacy Settings: A Tale of Opacity, Choice, and Discrimination. *PETS*, 2015.

[5] C. Dewey. 98 personal data points that Facebook uses to target ads to you. https://www.washingtonpost.com/news/the-intersect/wp/2016/08/19/98-personal-data-points-that-facebook-uses-to-target-ads-to-you.

[6] Facebook (Still) Letting Housing Advertisers Exclude Users by Race. https://www.propublica.org/article/facebook-advertising-discrimination-housing-race-sex-national-origin.

[7] Facebook Ads Preferences. https://www.facebook.com/ads/preferences.

[8] Facebook Advertising Policy: Personal Attributes. https://www.facebook.com/policies/ads/prohibited$_c$ontent/personal$_a$ttributes.

[9] Facebook Archive of Ads with Political Content. http://facebook.com/politicalcontentads/.

[10] Facebook Enabled Advertisers to Reach 'Jew Haters'. https://www.propublica.org/article/facebook-enabled-advertisers-to-reach-jew-haters.

[11] Facebook Leaks Usernames, User IDs, and Personal Details to Advertisers. http://www.benedelman.org/news/052010-1.html.

[12] Facebook Lets Advertisers Exclude Users by Race. https://www.propublica.org/article/facebook-lets-advertisers-exclude-users-by-race/.

[13] Facebook Messenger Chatbots Can Leak Your Private Information. https://www.techworm.net/2016/09/facebook-messenger-chatbots-can-leak-private-information.html.

[14] Facebook adds human reviewers after 'Jew haters' ad scandal. http://www.bbc.com/news/technology-41342642.

[15] Facebook: What are my ad preferences and how can I adjust them? https://www.facebook.com/help/247395082112892.

[16] Improving Enforcement and Promoting Diversity: Updates to Ads Policies and Tools. http://newsroom.fb.com/news/2017/02/improving-enforcement-and-promoting-diversity-updates-to-ads-policies-and-tools/.

[17] Google Ads Settings. https://adssettings.google.com/.

[18] Google Personalised Advertising Policy. https://support.google.com/adwordspolicy/answer/143465.

[19] Google: About audience targeting. https://support.google.com/adwords/answer/2497941?hl=en.

[20] How Google uses Customer Match data. https://support.google.com/adwords/answer/6334160.

[21] A. Korolova. Privacy Violations Using Microtargeted Ads: A Case Study. *Journal of Privacy and Confidentiality*, 3(1), 2011.

[22] B. Krishnamurthy, K. Naryshkin, and C. E. Wills. Privacy leakage vs. Protection measures: the growing disconnect. *IEEE W2SP*, 2011.

[23] B. Liu, A. Sheth, U. Weinsberg, J. Chandrashekar, and R. Govindan. AdReveal: Improving Transparency into Online Targeted Advertising. *HotNets*, 2013.

[24] M. Lecuyer, G. Ducoffe, F. Lan, A. Papancea, T. Petsios, R. Spahn, A. Chaintreau, and R. Geambasu. XRay: Enhancing the Web's Transparency with Differential Correlation. *USENIX Security*, 2014.

[25] M. Lecuyer, R. Spahn, Y. Spiliopolous, A. Chaintreau, R. Geambasu, and D. Hsu. Sunlight: Fine-grained Targeting Detection at Scale with Statistical Confidence. *CCS*, 2015.

[26] C. Patricia, C. Ruben, C. Angel, and K. Mikko. Independent Auditing of Online Display Advertising Campaigns. *HotNets*, 2016.

[27] J. Parra-Arnau, J. P. Achara, and C. Castelluccia. MyAdChoices: Bringing Transparency and Control to Online Advertising. *ACM TWEB*, 11, 2017.

[28] Protecting Privacy with Referrers. Facebook Engineering's Notes. http://www.facebook.com/notes/facebook-engineering/protecting-privacy-with-referrers/392382738919.

[29] T. Speicher, M. Ali, G. Venkatadri, F. N. Ribeiro, G. Arvanitakis, F. Benevenuto, K. P. Gummadi, P. Loiseau, and A. Mislove. On the Potential for Discrimination in Online Targeted Advertising. *FAT**, 2018.

[30] Sandvig V. Sessions — Challenge to CFAA Prohibition on Uncovering Racial Discrimination Online. https://www.aclu.org/cases/sandvig-v-sessions-challenge-cfaa-prohibition-uncovering-racial-discrimination-online.

[31] Shutting Down Partner Categories. https://newsroom.fb.com/news/h/shutting-down-partner-categories/.

[32] Tailored Audiences File Data. https://dev.twitter.com/ads/audiences/file-data.

[33] Twitter Ads Policies. https://business.twitter.com/en/help/ads-policies/introduction-to-twitter-ads/twitter-ads-policies.html.

[34] Twitter Ads Transparency Center. https://ads.twitter.com/transparency.

[35] G. Venkatadri, E. Lucherini, P. Sapiezyński, and A. Mislove. Investigating sources of PII used in Facebook's targeted advertising. *PETS*, 2019.

[36] G. Venkatadri, Y. Liu, A. Andreou, O. Goga, P. Loiseau, A. Mislove, and K. P. Gummadi. Privacy Risks with Facebook's PII-based Targeting: Auditing a Data Broker's Advertising Interface. *IEEE S&P*, 2018.

[37] C. E. Wills and C. Tatar. Understanding What They Do with What They Know. *WPES*, 2012.

[38] What Happens When I Upload My Customer List to Facebook? https://www.facebook.com/business/help/112061095610075.

[39] What's a Custom Audience from a Customer List? https://www.facebook.com/business/help/341425252616329/.

[40] Your Twitter data. https://twitter.com/settings/your$_t$witter$_d$ata.

[41] eyeWnder_Experiment. http://www.eyewnder.com/.