

The Accuracy of the Demographic Inferences Shown on Google’s Ad Settings

Michael Carl Tschantz Serge Egelman Jaeyoung Choi
Nicholas Weaver Gerald Friedland

International Computer Science Institute

October 3, 2016

Abstract

Google’s Ad Settings shows the gender and age that Google has inferred about a web user. We compare the inferred values to the self-reported values of 502 survey participants. We find that Google often does not show an inference, but when it does, it is typically correct. We explore which usage characteristics, such as using privacy enhancing technologies, are associated with Google’s accuracy.

1 Introduction

Google’s Ad Settings offers users a window into the model that Google learns about them from online tracking and their account settings [5]. Users may see inferences Google made about them at <https://www.google.com/settings/ads>. Figure 1 shows a screen shot of the first author’s settings page. Google designed this free service to provide transparency into their operations and to provide users with a chance to correct their data.

The page provides two predictions each for the user’s gender and age: one based upon the information Google uses for its web services and one based upon the information Google uses as part of its web-wide ad network. The page allows editing the inferences. Google provides some information about how it works [5], but questions remain about the accuracy of Google’s profiles.

To study the accuracy of Google’s predictions and how they are associated with user behaviors, attitudes, and usage of privacy enhancing technologies (PETs), we conducted a survey. We asked participants for their ages, genders, computer usage habits, attitudes, PETs usage, and for a copy-and-paste of the content of their Ad Settings page. We compared their supplied demographics to the age range and gender provided by Google and examined how various factors are associated with accuracy.



Figure 1: Screen shot of part of Google’s Ad Settings webpage while logged in with a Google account using Safari

We find that Google’s predictions tend to be accurate when Google makes them, but that Google often makes no inference. In particular, Google rarely makes predictions for logged out users. While we document that Google’s accuracy for some subgroups of users is far from its overall accuracy, we do not find statistical significance for any such association after adjusting for the large number of hypotheses examined in this exploratory work.

We believe we are the first look at the accuracy of Google’s inferences on Ad Settings with a survey. We provide a new point of reference for understanding Google’s ability to infer attributes of users. Additionally, we believe this paper is a reasonable starting point for larger-scale confirmation studies.

After covering related work, we discuss our research methods, and then present our results. We end with discussion. The appendix provides some additional details about the data we collected.

2 Related Work

Datta et al. experimented on Ad Settings to determine how they impact the ads shown and how browsing behaviors impact them [4]. Our work differs by looking at the accuracy of Google’s stated inferences on real users.

Balebako et al. studied the effectiveness of PETs by examining how personalized the ads shown to browsers with PETs were compared to those shown to browsers without PETs installed [1]. In addition to differing by looking at real users, our work differs by looking at the Ad Settings interface instead of ads.

Small-scale anecdotal examinations of the accuracy of Ad Settings have appeared in the popular press [2, 6], as has a survey looking at the accuracy of Google’s geo-location abilities [3].

3 Method

With IRB approval, we conducted a survey that consisted of three types of information collection. First, we provided participants standard questions to which they responded. We asked questions about their gender, age, browser usage, PETs usage, and opinions on the importance of privacy.

Second, we collected the inferred demographics displayed by Google’s Ad Settings to compare to the participant’s self-reported gender and age. We showed participants a screen shot of what this page looks like and included instructions on how to copy and paste the main content of the page. We asked them to paste it into a web form. We used scripts to extract various variables from this page: inferred age from data from Google’s services, inferred gender from Google’s services, inferred age from data from across the web (Google’s ad network), inferred gender from across the web, whether the user was logged into Google, and whether the user opted out of Google’s interest-based ads.

Third, we conducted measurements of the participants’ web browsers. We used an invisible iframe to have their browsers visit our server, which ran a series of tests to determine whether or not first party or third party cookies were blocked, whether Google Analytics cookies were blocked, and whether their browser was transmitting the DNT header.

Our survey yielded four measures of accuracy along with numerous factors they could be associated with, making a myriad of comparisons possible. To compensate for the multiple testing problem, we split our survey responses into an exploratory set and a confirmation set. We used the exploratory set to identify associations that appear statistically significant under the χ^2 test without adjusting for multiple tests. We then tested just these associations on the confirmation data while adjusting for the total number of confirmation tests (nine) using the χ^2 test with a Bonferroni correction. For reasons of space, we report frequencies and accuracies over the whole data set despite comments on statistical significance referring to the two subsets of data.

4 Results

We recruited for our survey using Mechanical Turk with an offer to pay 500 participants. On Oct. 29, 2014, 558 Turkers started our survey with 501 completing it. We eliminated 13 responses for not correctly providing us with a copy of their Ad Settings page and an additional 3 for not self-reporting gender and age.

Of the remaining responses, 4 of them were created using browsers with cookies disabled for which Google displayed a message saying as much and that it had no inferences for the person. To simplify the presentation, we eliminated these responses as well, although they could be considered additional cases of Google opting to not make an inference.

We took the the first 289 of the remaining 481 responses to be our exploratory data set and the remaining 192 responses to be our confirmation data set.

Account Settings. Table 1 shows how many of the remaining respondents opted out of tracking on Google services or Google ads across the web. Additionally, we found that 13 of the respondents who opted out of Google ads on services also opted in for getting ads on YouTube, presumably overriding the more general opt out for that service.

	Google	Across the web
not opted out	455	467
opted out	26	14

Table 1: Number of respondents who opted out or in of various forms of tracking by Google

	Reported		Google inferred on			
	in	out	Services		Web	
			in	out	in	out
Female	159	35	102	0	130	11
Male	238	49	178	0	195	11
Unknown	n/a	n/a	78	0	48	40
N/A	n/a	n/a	39	84	24	22
18-24	69	8	49	0	45	2
21-35	n/a	n/a	1	0	1	0
25-34	186	41	144	0	148	8
35-44	92	17	70	0	80	2
45-54	21	11	21	0	28	2
55-64	20	5	9	0	17	2
65+	9	2	6	0	6	2
Unknown	n/a	n/a	58	0	48	44
N/A	n/a	n/a	39	84	24	22

Table 2: Number of respondents with each value for Google’s inferred gender and age broken by being logged in or out

We found 397 respondents to be logged into an Google account and 84 to not be. Since logged users’ profiles are available, which makes inferences easier, we break down all further results along the lines of logged in and out users.

Demographics. Table 2 shows both the self-reported and inferred genders and ages. The values of “Unknown” and “N/A” are ones that Google lists, not comments from the authors on what we know or applicability. (“n/a” is ours.) Google gave one user an age range that overlapped with two other age ranges. We drop this range from further analysis.

Table 3 summarizes Google’s accuracy obtained from the responses. They show that when Google makes a prediction, it is typically correct. However, Google often avoids making a prediction. In particular, Google never makes predictions for logged out users for Google services and rarely does so for logged out users for ads across the web.

We checked whether the percentage Google got right was associated with gender or age. The results, shown in Table 4, include two statistically significant

	Right		Wrong		Skipped	
	in	out	in	out	in	out
Sex Google services	66	0	5	0	29	100
Sex Across the web	74	21	8	5	18	74
Age Google services	67	0	9	0	24	100
Age Across the web	65	11	16	11	18	79

Table 3: Google’s accuracy shown in percentages. We treat as skipped those Google called “Unknown” or “N/A”.

	Sex services		Sex web		Age services		Age web	
	in	out	in	out	in	out	in	out
All	66	0	74	21	67	0	65	11
Female	61	0	75	26	63	0	69	17
Male	69	0	73	18	70	0	63	6
18-24	71	0	74	38	68	0	57	25
25-34	66	0	75	22	71	0	69	12
35-44	62	0	72	12	68	0	72	6
45-54	71	0	76	27	76	0	81	9
55-64	55	0	65	20	35	0	45	0
65+	78	0	78	0	11	0	11	0

Table 4: The percentage that Google got right for each reported gender and age range. Boldface shows associations with statistical significance in our exploratory data set.

associations for age in the exploratory data set. However, neither of these differences proved significant in our confirmation data set.

Computer Usage. Table 5 shows the number of respondents with various usage conditions on the computer used to take our survey and Google’s rate of making correct inferences for them. Some of these activities intuitively makes it more difficult to correctly make inferences about any one user of the computer since they imply that the computer has multiple users, which could pollute a model of any one of them. The decrease in accuracy is sizable in some cases and reaches statistical significance in our exploratory data set for three conditions involving the clearing cookies. However, none have a statistically significant association with Google’s error rate in our confirmation data set.

Attitudes. Table 6 shows the associations between the respondents’ attitudes toward tracking and Google’s accuracy. Our exploratory analysis found no significant associations.

PETs. Table 7 shows the usage of various PETs. Above the bar are the self-

	Counts				Percent right							
	Yes		No		Sex services		Sex web		Age services		Age web	
	in	out	in	out	in	out	in	out	in	out	in	out
All (baseline)	n/a	n/a	n/a	n/a	66	0	74	21	67	0	65	11
Shared computer	62	19	335	65	63	0	77	21	71	0	69	11
Shared account	28	12	369	72	54	0	75	25	64	0	61	17
Other users in a week	107	24	290	60	62	0	74	21	69	0	69	17
Other users yesterday	61	17	336	67	59	0	67	18	66	0	64	18
Cleared cookies today/yesterday	58	36	339	48	64	0	66	8	64	0	59	3
Cleared cookies on close	33	17	364	67	61	0	64	18	55	0	48	6
Private mode	18	6	379	78	56	0	61	0	61	0	56	0

Table 5: The number of respondents with each computer usage characterization and the percentage that Google got right for respondents with each computer usage characterization. For the percentages, boldface shows associations with statistical significance in our exploratory data set (i.e., a statistically significant difference from the baseline). We defined a “shared computer” to be one that respondent described as “Regularly used by multiple workers at a place of employment”, “Regularly used by multiple members of a family”, or “Regularly used by many people in a public place (library, Internet cafe, etc.)”, but not as “Regularly used only by me” nor as “None of the above”.

	Sex services	Sex web	Age services	Age web
	in	out	in	out
All (baseline)	66	0	74	21
Concerned about tracking	62	0	71	25
Confidence about avoiding it	62	0	73	23
More concerned about accurate profiles	69	0	77	28
More concerned about inaccurate profiles	72	0	77	14

Table 6: The percentage that Google got right for respondents with each attitude. We defined “Concerned about tracking” as answering with a 4 or 5 (very concerned) on a 5-point scale to the question “How concerned are you about online tracking of your behavior?” We defined “Confidence about avoiding it” as a 4 or 5 (very confident) to the question “If you have taken steps to prevent online tracking of your behavior, how confident are you that it prevents online tracking?”

	Yes		No		I don't know		empty	
	in	out	in	out	in	out	in	out
AdBlock	216	45	138	28	40	11	3	0
Ghostery	18	8	333	69	45	7	1	0
NoScript	12	10	317	54	66	19	2	1
DoubleClick opt out	15	4	298	66	84	14	0	0
Webpage opt out	82	20	246	47	69	16	0	1
DNT set	100	34	217	33	80	17	0	0
DNT sent	53	22	339	55	0	0	5	7
1st cookies off	22	11	370	66	0	0	5	7
3rd cookies off	24	13	367	64	0	0	6	7
Google cookies off	39	17	353	60	0	0	5	7

Table 7: Number of users of each PET

reported usage habits of PETs by respondents. Below the bar are our server’s measurements. For these measurements, “empty” means that our server did not detect a visit from the respondent (e.g., due to network loss). Table 8 shows Google’s accuracy for users of each PET.

Looking at AdBlock, in our exploratory data set, we found a statistically significant reduction in the accuracy of Google for data from across the web for gender both when logged in and out. A significant reduction in the accuracy for age across the web also exists, but only when logged in.

The only other PET to get statistically significance is using webpage opt outs, and only in the case of age across the web when logged in.

None of these differences proved statistically significant in our confirmation data set.

5 Discussion

We have no way of knowing whether the inferences shown on Ad Settings are the same as those actually used by Google for ad targeting, and prior work suggests that Ad Settings does not provide information about how Google remarkets to users based upon prior webpage visits [4] (a limitation made explicit on the Ad Settings page after the publication of [4]). Nevertheless, we find it noteworthy that Google rarely shows inferences for logged out users. We can only conjecture as to the reason, but perhaps one’s web browsing behavior is not as visible to or interpretable by Google as some fear. Unfortunately, since conducting our survey, Google has disabled Ad Settings for logged out users, precluding the possibility of further studying this phenomenon.

We relied upon self-reports of age and gender for ground truth, of PETs usage, and of behavior while looking for factors associated with Google’s accuracy. Self-reports of PETs usage, in particular, may be inaccurate due to

	Sex services		Sex web		Age services		Age web	
	in	out	in	out	in	out	in	out
All (baseline)	66	0	74	21	67	0	65	11
AdBlock	66	0	67	7	65	0	56	4
Ghostery	56	0	33	0	61	0	28	0
NoScript	42	0	75	0	58	0	58	0
DoubleClick opt out	53	0	60	50	47	0	60	25
Webpage opt out	65	0	61	30	60	0	57	15
DNT set	57	0	68	18	65	0	65	6
DNT sent	57	0	68	23	62	0	57	9
1st cookies off	45	0	68	36	68	0	68	9
3rd cookies off	50	0	71	38	71	0	75	15
Google cookies off	54	0	64	24	59	0	59	6

Table 8: The percentage that Google got right for users of each PET. Boldface shows associations with statistical significance in our exploratory data set.

the obscurity of PETs and the possibility that a shared browser may use one without the respondent’s knowledge. Furthermore, our server’s attempts to detect PETs usage by examining the behavior of respondents’ browsers could have measurement errors from factors such as network loss.

Mechanical Turkers might not be representative of standard web users. In particular, they may be more likely to use PETs or security measures due to the heavy use of their browsers for Turking. Furthermore, they may visit an atypically large number of webpages unassociated with their demographics to fulfill their Turking tasks.

Our exploratory results suggest that cookie clearing and AdBlock may be associated with decreasing Google’s accuracy. Using observational data, we cannot claim that they cause the decrease.

Future work includes running experiments to determine whether PETs usage is the cause of such decreases in accuracy and conducting a larger-scale observational studies to bring a larger number of PETs users and cookie clearers into the sample. We hope this will allow us to find statistically significant associations, which in some cases appear unobtainable due to the small number of respondents with privacy-seeking behaviors (Tables 5 and 7). However, we also worry that our preliminary results may suggest that some PETs, under typical use by real users, have little impact on the accuracy of Google’s inferences.

Acknowledgements

This research was supported by the U.S. National Science Foundation (NSF) grants CNS 1065240 and CNS 1514509. The views and conclusions contained in this document are those of the authors and should not be interpreted as

representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

References

- [1] BALEBAKO, R., LEON, P., SHAY, R., UR, B., WANG, Y., AND CRANOR, L. Measuring the effectiveness of privacy tools for limiting behavioral advertising. In *Web 2.0 Security and Privacy Workshop* (2012).
- [2] BOSCH, T. Does Google accurately guess your age and gender? *Slate* (Jan. 2012). www.slate.com/blogs/future_tense/2012/01/25/google_ad_preferences_manager_does_it_accurately_guess_your_age_and_gender_.html.
- [3] CAMERON, S. Does Google really know where you are? for nearly half of you, the answer is no. *Search Engine Land* (2015). searchengineland.com/google-really-know-230001.
- [4] DATTA, A., TSCHANTZ, M. C., AND DATTA, A. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. In *Privacy Enhancing Technologies (PoPETs)* (2015), pp. 92–112.
- [5] GOOGLE. About Ads Settings. Ads Help webpage: <https://support.google.com/ads/answer/2662856?hl=en>, 2016.
- [6] LARSON, S. I crawled my ad settings to see what Google really knows about me. *The Daily Dot* (2015). www.dailydot.com/debug/google-data-ad-settings/.

A Additional Data

The following tables show the responses we received to various questions on our survey. Figure 2 provides a larger screen shot of Google’s Ad Settings.

	Right		Wrong		Skipped	
	in	out	in	out	in	out
Sex Google services	61	0	9	0	30	100
Sex Across the web	75	26	13	6	12	69
Age Google services	63	0	6	0	31	100
Age Across the web	69	17	18	11	13	71

Table 9: Google’s accuracy shown in percentages for females.

Ads Settings

Settings for Google ads

Ads enable free web services and content. These settings help control the types of Google ads you see.

	Ads on Google	Google ads across the web [?]
	 Search  Gmail  YouTube  Maps	 Google ads across the web
Gender	Unknown Visit your Google Profile	Male Edit Based on the websites you've visited
Age	25-34 Visit your Google Profile	25-34 Edit Based on the websites you've visited
Languages	N/A	English Edit Based on the websites you've visited
Interests	Unknown Edit From your previous activity on Google	Air Travel, and 15 more Edit Based on the websites you've visited
Advertisers' campaigns you've blocked [?]	None From your blocking activity	N/A
Opt-out settings	Opt out of interest-based ads on Google	Opt out of interest-based Google ads across the web

Visit the [Consumer Ads Help Center](#) to learn more about how Google serves ads.

Google adheres to advertising industry privacy standards. To learn about these standards, including how you can opt out of interest-based advertising from Google and other participating companies, visit our [About Google Ads](#) page. If you want to permanently opt out of the DoubleClick cookie, you can install the [DoubleClick opt out extension](#).

[Send feedback](#)

Figure 2: Screen shot of Google's Ad Settings webpage while logged in with a Google account using Safari

	Right		Wrong		Skipped	
	in	out	in	out	in	out
Sex Google services	69	0	2	0	29	100
Sex Across the web	73	18	5	4	22	78
Age Google services	70	0	10	0	20	100
Age Across the web	63	6	15	10	22	84

Table 10: Google’s accuracy shown in percentages for males.

	in	out
Yes	389	80
I don’t know	1	0
empty	2	0
No	5	4

Table 11: Did you use this computer yesterday?

	in	out
5 days	70	11
6 days	17	3
7 (every day)	289	63
I don’t know	2	0
4 or fewer days	16	7
empty	3	0

Table 12: In the past week, on how many days did you use this computer?

	in	out
Regularly used only by me	331	65
Regularly used by multiple workers at a place of employment	6	1
Regularly used by multiple members of a family	53	17
Regularly used by many people in a public place (library, Internet cafe, etc.)	3	1
None of the above	4	0

Table 13: Which best describes this computer?

	in	out
3 or more days	57	14
0 days (no one else used it)	279	57
I don't know	11	3
2 days	18	5
1 day	32	5

Table 14: In the past week, on how many days did someone other than you use this computer?

	in	out
I don't know	12	3
No	124	22
Not applicable	218	45
Yes	28	12
Some	13	2
empty	2	0

Table 15: If anyone else used the computer you are currently using within the last week, did that person(s) use a different user account from the one you are currently using?

	in	out
Yes	61	17
I don't know	5	1
empty	0	1
No	331	65

Table 16: Did anyone else use this computer yesterday?

	in	out
1	62	9
0	123	21
3	86	25
2	71	15
5	4	1
4	51	12
empty	0	1

Table 17: If you have taken steps to prevent online tracking of your behavior, how confident are you that it prevents online tracking?

	in	out
inaccurate profile about you	53	7
accurate profile about you	134	25
empty	1	1
equally concerning	209	51

Table 18: Which is more concerning to you?

	in	out
1	14	3
3	95	15
2	44	7
5	51	21
4	190	38
empty	3	0

Table 19: How concerned are you about online tracking of your behavior?

	in	out
Week	83	11
Month+	96	10
I don't know	42	9
Never	56	11
Yesterday	41	19
Today	17	17
Month	62	7

Table 20: When was the last time you cleared the cookies of the browser you are currently using?

	in	out
Yes	18	6
I don't know	18	6
empty	1	1
No	360	71

Table 21: Are you currently using your web browser in "private browsing mode" (sometimes called "incognito")?

	in	out
Yes	33	17
I don't know	89	11
empty	5	2
No	270	54

Table 22: Does the browser you are currently using automatically clear cookies upon closing it?

	in	out
Some high school	4	2
Some college or Associate degree	135	32
Some graduate school	11	1
Master's degree	31	9
Bachelor's degree	147	23
Doctorate or professional graduate degree (Ph.D., J.D., M.D., etc.)	10	3
High school diploma or GED	59	14

Table 23: What is your highest completed level of education?