

How does Google know my gender if I didn't say it? Measuring how Google infers the gender of the users*

Pelayo Vallina^{1,2}, Antonio Fernández Anta¹, Rubén Cuevas², Angel Cuevas²

¹IMDEA Networks Institute

²Universidad Carlos III de Madrid

pelayo.vallina@imdea.org, antonio.fernandez@imdea.org, rcuevas@it.uc3m.es

Introduction

Nowadays, online services like social networks or e-commerce sites (among many others) have the capacity of learning the preferences and interests of individual users based on their online activity. Furthermore, some of these online services have also access to the users' personal data, like their phone number, their gender, their age, their neighborhood, or the places they visit. Using such knowledge and algorithms that they keep private, online services can generate very accurate user profiles, and use them to deliver personalized content. This personalization can be beneficial for the users, for example, when they receive good recommendations about films and TV series on their favorite streaming platform. However, personalization does not always benefit users. Previous work has shown that the algorithms used by online services to decide which content each particular user receives can in fact discriminate the user based on personal data, like gender, race, or geographical location (Hannak et al. 2014; Kliman-Silver et al. 2015; Hannák et al. 2017; Mikians et al. 2012). In addition, recent reports have revealed how Cambridge Analytica, a political consulting company, used Facebook data to try to influence users in national elections (Guardian 2018). All these facts have raised concerns in society, putting in the spotlight the importance of knowing which user data is collected by online services, how it is processed, and for which purposes.

Due to the demands of the online users and the pressure of new digital regulation (EU-GDPR), more online services have started to implement transparency tools like Google Ads Settings¹. This platform that belongs to Google, gives users the capacity to control which categories and interest have associated with them based on their behavior. Google computes such categories using the data that collects through different mechanisms, like the queries that users make using the search engine, the videos that they watch on YouTube or also thought tracking techniques (Ex: Cookies delivered thought ads). Also, Google Ads Settings gives

*Partially supported by la Caixa Foundation grant MyBubble (LCF/PR/MIT17/11820009), the Spanish MICIU grant DiscoEdge (TIN2017-88749-R), the Comunidad de Madrid grant EdgeData-CM (P2018/TCS-4499), and the NSF of China grant 61520106005. Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.google.com/settings/ads>.

users the capacity to editing the interests and categories associated with their account, as well as remove the historical data generated by their activities, as the videos watched on YouTube, the webs visited, and so on and so forth. Furthermore, Google Ads Settings not only provides the categories and interests associated with each account but also provides as a category the gender of the users even when these users don't provide such personal data. This makes us wonder how Google infers the users' gender.

- Is it a systematic phenomenon? or is it an isolated case what we have seen?
- If it is a systematic phenomenon, How does Google infer the gender of the users? Does it do it based on pure data analysis, for example, counting the number of females/males that have similar behavior?
- Could be possible that Google infers the gender based on categories and stereotypes? If Google does it, do the results change based on the user's location?

Methodology

We develop a methodology that will help us to answer the previous questions. This methodology follows some steps:

Creating Google Accounts

The first step is to create a representative number of Google accounts without historical data associated with them. However, we need to take some considerations when we create such accounts:

- We don't specify the gender of the account in the log up.
- The names of the accounts follow a specific pattern, `experimentXpersonaX`, where X is a random number. We want to create accounts with names that Google can't use to infer the gender.
- All the accounts have the same age and nationality, in our case, Spanish and 40 years old.

Personas

The second step in the methodology is to create synthetic profiles, *personas*, that pretend to be real users with a strong preference for a certain topic or a combination of them (Carrascosa et al. 2015). Each *persona* consists of

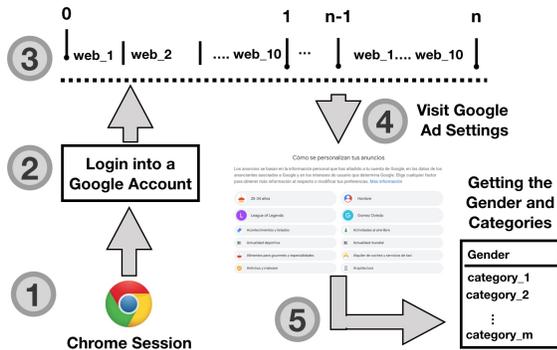


Figure 1: Process.

a set of 10 websites carefully selected and specialized in a certain category. We validated the categories of each websites using third-party services like McAfee², OpenDNS³, and VirusTotal⁴. As an example, *persona* “Football” consists in a set of ten websites specialized in content related with “Football”. Furthermore, each website should contain third-party services that belong to Google embedded in the website, like `google-analytics.com`, `doubleclick.com` among others, to be sure that Google monitors our online activity.

Providing behavior to our Google Accounts

Once we have created the Google accounts and the set of *personas*, it’s the moment to provide a specific behavior to our Google accounts using those *personas*. To achieve that aim, we follow a methodology that consists of 5 different steps, as we present in Figure 1. For each new Google account that we have created, we open a new 1) Chrome session without historical data using selenium webdriver⁵. Then 2) we login into the Google account. To give a strong behavior to our account, we choose a *persona* from our set and 3) we visit the 10 websites that compose such *persona* “N” number of times using the Chrome browser. We wait 8 hours in order to give time to Google to calculate and associate the categories to our Google account. Finally, after that time, 4), we visit Google AdSettings to 5) collect the categories and the gender associated with the account. In case that there is no gender, we compute it as a *no result*. We repeat the same experimentation with all the Google accounts that we have used a different *personas*.

Preliminary Results

In this work in progress, our first objective is to measure how often Google infers the gender of the users that haven’t provide such information. We use 36 different *personas* associated with a different Google account. As we can see in Table 1, at the end of the experiment, 28 accounts were

associated to Male gender, while 8 accounts were associated to Female gender and `no_result`, concluding that it is a systematic phenomenon. Also, we found some interesting patterns, highlighting the association of “Soccer” and “Business” with Male and “Pregnancy” and “Cuisine” with Female.

Table 1: Results of the 36 experiments. This table shows the relationship between the *persona* and the gender that Google infers, which can be Male, Female or `no_result` in case that Google wasn’t able to compute the gender of the user.

Persona	Gender	Persona	Gender	Persona	Gender
7-12Education	Female	AirTravel	Male	Cuisine	Female
Antiques	Male	ArcadeCoinGames	Male	Automotive	Male
Beauty	Female	Business	Male	Clothing	<code>no_result</code>
Construction	Female	Education	<code>no_result</code>	FinancialNews	Male
Guitar	Male	HobbiesInterests	<code>no_result</code>	HomeGarden	Female
HVACClimate	<code>no_result</code>	InternetTec.	<code>no_result</code>	JobFairs	Male
KitchenDining	<code>no_result</code>	LegalIssues	Male	Motorcycles	Male
Movies	Male	Music	Male	News	<code>no_result</code>
PersonalFinance	Male	Pregnancy	Female	Basketball	Male
RealEstate	Male	Running-Jogging	Male	Science	Male
Shopping	Female	Sports	Male	Technology	<code>no_result</code>
Travel	Female	ComputerGames	Male	Soccer	Male

Conclusions and Future Work

We have detected that Google infers the gender of the users based on its online behavior even when users don’t provide it. However, we were not able to verify how Google infers users’ gender. We plan to repeat and extend the experiments in order to have a better and more complete knowledge of the process. Furthermore, we will run experiments changing the location of the users to study if the location affects the results.

References

Carrascosa, J. M.; Mikians, J.; Cuevas, R.; Erramilli, V.; and Laoutaris, N. 2015. I always feel like somebody’s watching me: measuring online behavioural advertising. In *CoNext’15*.

EU-GDPR. GDPR . <https://gdpr-info.eu/>.

Guardian, T. 2018. How Cambridge Analytica turned Facebook likes into a lucrative political tool . <https://www.theguardian.com/technology/2018/mar/17/facebook-cambridge-analytica-kogan-data-algorithm>.

Hannak, A.; Soeller, G.; Lazer, D.; Mislove, A.; and Wilson, C. 2014. Measuring price discrimination and steering on e-commerce web sites. In *IMC’14*, 305–318.

Hannák, A.; Wagner, C.; Garcia, D.; Mislove, A.; Strohmaier, M.; and Wilson, C. 2017. Bias in online freelance marketplaces: Evidence from taskrabbit and fiverr. In *ACM CSCW’17*.

Kliman-Silver, C.; Hannak, A.; Lazer, D.; Wilson, C.; and Mislove, A. 2015. Location, location, location: The impact of geolocation on web search personalization. In *IMC’15*. ACM.

Mikians, J.; Gyarmati, L.; Erramilli, V.; and Laoutaris, N. 2012. Detecting price and search discrimination on the internet. In *HotNets’12*.

²<https://www.mcafee.com/enterprise/en-us/threat-center.html>

³<https://community.opendns.com/domaintagging/>

⁴<https://www.virustotal.com/>

⁵<https://www.seleniumhq.org/projects/webdriver/>