



MAX PLANCK INSTITUTE
FOR SOFTWARE SYSTEMS



Northeastern

Inria

Investigating ad transparency mechanisms in social media

Oana Goga
CNRS, Univ. Grenoble Alpes

Work done with Athanasios Andreou, Giridhari Venkatadri,
Krishna P. Gummadi, Patrick Loiseau, Alan Mislove



In this talk

Explanations for social media targeted advertising

Ad explanations

Data explanations



Why am I being shown this ad?

What data the ad platform knows about me?



[NDSS'18]

Facebook provides explanations

The screenshot shows a Facebook advertisement for TransferWise. The ad text reads: "Sending money to India? Don't get caught out by bad exchange rates. Save time and money using TransferWise." To the right of the ad is a menu with three options: "Save video" (Add this to your saved items), "Hide ad" (See fewer ads like this), and "Report ad" (Mark ad as offensive or inappropriate). Below the ad, a section titled "Why Am I Seeing This Ad?" is visible. It contains two paragraphs of text explaining the targeting: "One reason you're seeing this ad is that Leibniz wants to reach people interested in Ferrero SpA, based on activity such as liking Pages or clicking on ads." and "There may be other reasons you're seeing this ad, including that Leibniz wants to reach people ages 25 and older who live or were recently in Germany. This is information based on your Facebook profile and where you've connected to the internet." Below the explanation is a video player with the title "I have absolutely no clue" and the text "Stop using your bank, they are charging you a lot of money!". A "Learn More" button is located at the bottom right of the video player. The URL "TRANSFERWISE.COM" is visible at the bottom left of the video player.

TransferWise
Sponsored · 🌐

Sending money to India? Don't get caught out by bad exchange rates. Save time and money using TransferWise.

Save video
Add this to your saved items

Hide ad
See fewer ads like this

Report ad
Mark ad as offensive or inappropriate

Why Am I Seeing This Ad?

Options ▾

One reason you're seeing this ad is that Leibniz wants to reach people interested in Ferrero SpA, based on activity such as liking Pages or clicking on ads.

There may be other reasons you're seeing this ad, including that Leibniz wants to reach people ages 25 and older who live or were recently in Germany. This is information based on your Facebook profile and where you've connected to the internet.

I have absolutely no clue

"Stop using your bank, they are charging you a lot of money!"

TRANSFERWISE.COM

Learn More

Explications are voluntary or to satisfy law

But explanations are not trivial

- The **systems** they have to explain are **complex**
- Many design choices
 - Format, length, amount of details ...
- What is a **good explanation**?
 - Improve **control**
 - Satisfies **curiosity**
 - Detect **malicious or deceiving advertising**
 - Verify **compliance**

Why did I received this ad?

Ad explanations

... it's complicated

Targeted advertising is a complex system

- Facebook inferred some attributes
- Advertiser used attributes to select audience
- Facebook matched the ad to me through auctions

Desired properties of explanations

- Do explanations show all the attributes? (**completeness**)
- Were the attributes showed actually used by the advertisers? (**correctness**)
- Are explanations specific to each user? (**personalization**)
- Are explanations consistent across time? (**consistency**)

We need standards for explanations

To protect against **adversarial explanations** :

- **Insufficient / unsatisfactory**
 - That offer no insightful/actionable information to consumers
- **Misleading / fake**
 - Designed to gain consumer acceptance for a service
 - Misled consumers about the process

Measurement methodology

- Chrome extension to **collect ads from Facebook timeline**
 - 35 users for 5 months
 - 26K unique ads and explanations
 - **Controlled experiments** targeting users with ads:
 - We targeted users
 - We collected explanations
 - 96 successful campaigns
- } **Ground truth**

Are Facebook explanations complete?

- For ads targeting **customer PII**
 - “One reason you're seeing this ad is that [Booking.com](#) added you to a list of people they want to reach on Facebook. They were able to reach you because you're on their customer list or you've provided them with your contact information off of Facebook. This is based on [customer information](#) provided by Booking.com..”
- Does not show **what PII** booking.com used!
 - Email ? Telephone ? Name+address? etc.

Are Facebook explanations complete?

- For ads targeting **data broker attributes**
 - “One reason you're seeing this ad is that **Peugeot** wants to reach people who are part of an audience created based on **data provided by Acxiom**. Facebook works with data providers to help businesses find the right audiences for their ads. Learn more about data providers.”
- Does not say **what Acxiom provided attributes** were used!
 - Financial data ? Purchasing habits ? etc.

Are Facebook explanations complete?

- For ads targeting **Facebook attributes**:
 - “One reason you're seeing this ad is that **Peek & Cloppenburg** wants to reach people interested in **Shopping and fashion**, based on activity such as liking Pages or clicking on ads.”
 - “There may be other reasons why you're seeing this advert, including that Acer wants to reach **people aged 18 to 45 who live or have recently been in Germany**. This is information based on your Facebook profile and where you've connected to the Internet.”
- Picks exactly **one attribute** (besides gender, location, age)

Validation of incompleteness

- Ran several **controlled ads** targeting ourselves using a custom list and selecting **millennial & expats**
 - “One of the reasons why you're seeing this advert is because we think that you may be in the **Millennials** audience. This is based on what you do on Facebook..”
- Only one features, **millennial (not expats)**, shown!

Do explanations need to be complete?

- Should they specify **all attributes in ad targeting?**
- Arguments **for**:
 - Avoid misleading and insufficient explanations:
 - Designed to gain consumer acceptance for a service
 - Builds trust and incentivizes cooperation
- Arguments **against**:
 - Targeting formula may be a business secret
 - Overloads users with information (need succinct explanations)

Selecting attributes for explanations

“One reason you're seeing this ad is that [Peek & Cloppenburg](#) wants to reach people interested in [Shopping and fashion](#), based on activity such as liking Pages or clicking on ads.”

- Are the explained attributes **the most important**?
- Is [Shopping and fashion](#) the most important of all the [user' attributes](#) that Facebook and the advertised used to target the user?

How Facebook selects attributes

- Ran **controlled ads** to **reverse-engineer** Facebook's feature selection strategy
- Facebook appears to **prioritize attributes** based on
 - Their **type**: Demographic > Interest > PII > Behavioral
 - Their **prevalence**: Most prevalent first
- Unclear, if this is the **right prioritization** for users

Are Facebook explanations (at least) correct?

- Experiment : Ran a **controlled ad** targeting ourselves using a **custom list** and selecting **millennials & expats**
 - “**There may be other reasons** why you're seeing this advert, including that Vacations in Saarbücken wants to reach **people aged 18 and above who live or have recently been in Germany**. This is information based on your Facebook profile and where you've connected to the Internet.”
- **Never used attributes shown in explanations!**

Need for rigorous explanations

Incomplete explanations:

- Malicious advertiser can conceal sensitive/discriminatory attributes by adding a common popular attribute to the targeting audience

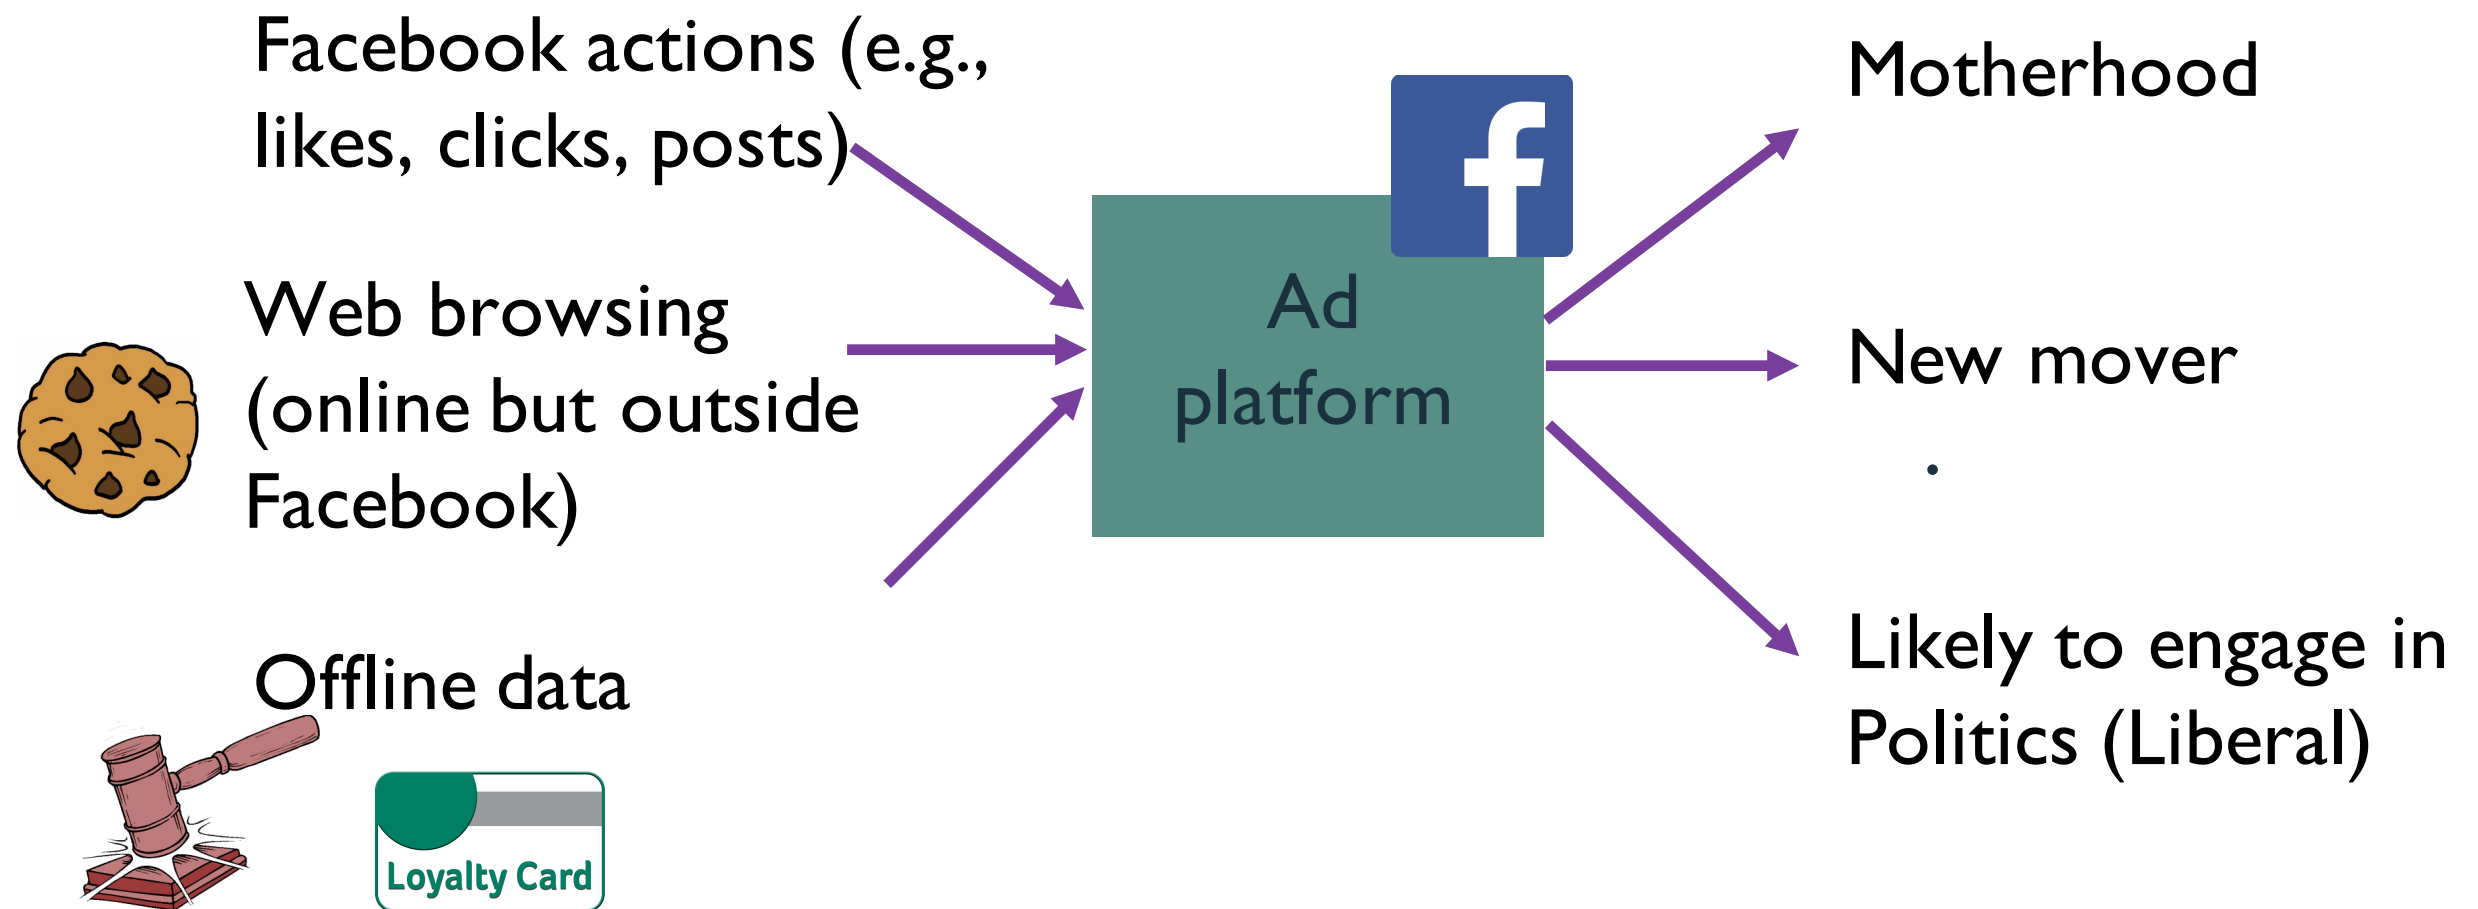
Misleading explanations:

- Fail to capture accurately the reasons why a user is targeted → induce false sense of trust


What data the ad platform knows about me?

Data explanations

How is the data inferred?



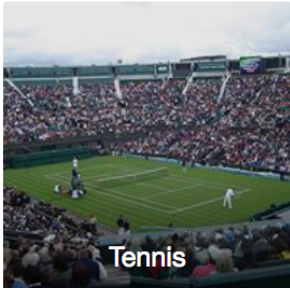
Explanations of the data inference process

 **Your interests** Close ^

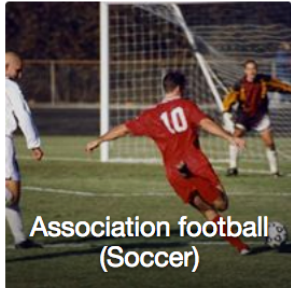
About Tennis
You have this preference because you liked a Page related to Tennis.

People Travel, places and events Technology **Sports and outdoors** More

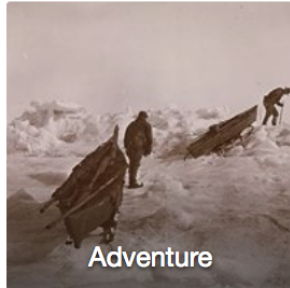
on Facebook or remove it from your ad preferences.



Tennis



Association football (Soccer)



Adventure

 **Your information** Close ^

About you **Your categories**

The categories in this section help advertisers reach people who are most likely to be interested in their products, services, and causes. We've added you to these categories based on information you've provided on Facebook and other activity.

Birthday in 02 February	Close Friends of Women with a Birthday in 0-7 days
Away from hometown	Close Friends of People with a Birthday in 0-7 days
Life, Physical, and Social Science	Architecture and Engineering
Healthcare and Medical	Smartphone Owners
Primary Browser: Chrome	4G Connection
All Android devices	Tablet Owners

Ad Preferences Page

Desired properties of a data explanation

- Specificity
- Completeness
- Correctness

Measurement methodology

- Build tool that collects the Ad Preference Page daily
 - ▶ Collection of real-world data
 - ▶ Controlled ad campaigns

Evaluation of properties

- Most explanations are **vague**

Pattern	No explanations
You have this preference because you liked a Page related to <i>interest</i>	2,866
You have this preference because you clicked on an ad related to <i>interest</i>	2,095
You have this preference because we think it may be relevant to you based on what you do on Facebook, such as pages you've liked or ads you've clicked	785
You have this preference because you installed the app <i>app – name</i>	107
This is a preference you added.	2

- Explanations are **incomplete**
 - No data broker attributes appear

Need for rigorous explanations

Incomplete explanations:

- Does not show the full picture to the user
- Provides a false sense of trust

Vague explanations:

- Does not allow users to control the outputs in the future

Takeaways

- Just mandating explanations is not enough!
- Badly designed explanations can be dangerous
 - Easily exploitable by malicious advertisers
- Designing good explanations is complicated
 - Different purposes ask for different properties

Open challenges

- How to pick **a few (K) features** for explanations?
- How to determine the **importance of a user attribute**?
 - Does it reveal **privacy sensitive information**?
 - Is it a **rare (or low prevalence) attribute in population**?
 - Does it exert **the most influence**?
- What properties explanations need to protect against malicious advertisers?

**A step towards more
transparency**

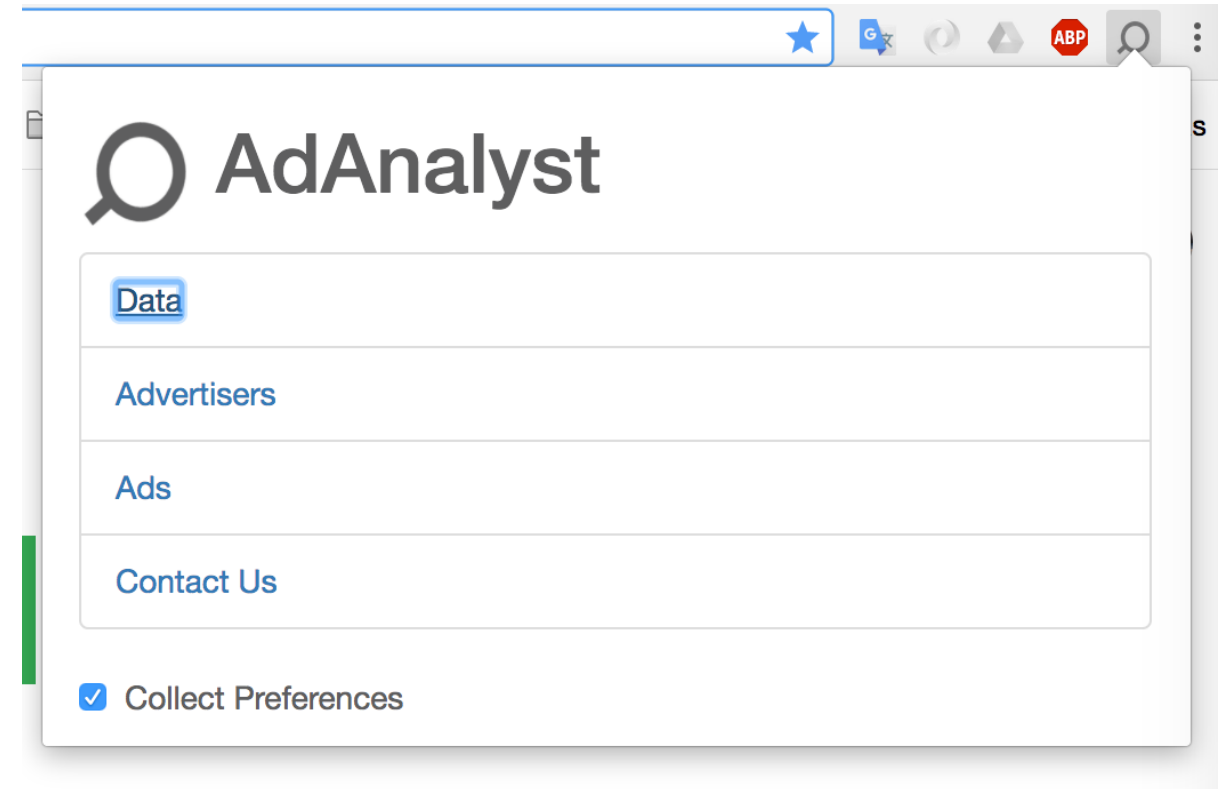

AdAnalyst

Make sense of the ads you receive on Facebook

- Enhance transparency by aggregated statistics
- Enhance transparency in a collaborative way

<http://adanalyst.mpi-sws.org/>

Disable/pause AdBlockPlus on Facebook!



Ads view

Ads

Check your daily unique ad consumption, by clicking on any date in the calendar!

December 2017						
Su	Mo	Tu	We	Th	Fr	Sa
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31						



Download this machine learning eBook to learn basic concepts and explore advanced algorithms.

Time: 11/12/2017 00:11

Advertiser: [MATLAB](#)

Advertiser category: Software

Explanation: One reason you're seeing this ad is that MATLAB wants to reach people who may be similar to their customers. Learn more. There may be other reasons you're seeing this ad, including that MATLAB wants to reach people ages 25 to 55 who live in France. This is information based on your Facebook profile and where you've connected to the internet.

Compact explanation: Lookalike Audience

Users also received this ad for: Lookalike Audience

The advertiser targeted other users with: Lookalike Audience, Computer vision (Interests), Data science (Interests), Machine learning (Interests)



Wish www.wish.com What other products are on sale under 5€ today?

Time: 11/12/2017 00:09

Advertiser: [Wish](#)

Advertiser category: App Page

Explanation: You're seeing this ad because Wish wants to reach women ages 16 to 60 who live in Germany. This is information based on your Facebook profile and where you've connected to the internet.

Compact explanation: 16-60 (Age/Gender/Location)

The advertiser targeted other users with: English (US) (Biographical Data)

Data view

Latest data Facebook has inferred about you

Latest interests:



Diane Kruger

People

You have this preference because you clicked on an ad related to Diane Kruger.

Added on: 08/12/2017




Term of endearment

Business and industry

You have this preference because you clicked on an ad related to Term of endearment.

Added on: 04/12/2017



Photograph

Hobbies and activities

You have this preference because you clicked on an ad related to Photograph.

Added on: 04/12/2017



Telephone

Business and industry

You have this preference because you clicked on an ad related to Telephone.

Added on: 04/12/2017



Short Message Service

Technology

You have this preference because you clicked on an ad related to Short Message Service.

Added on: 03/12/2017

Latest demographics:

Anniversary within 31-60 Days

People with a relationship anniversary (marriage, domestic partnership, etc.) occurring wi...

Added on: 08/12/2017

Life, Physical and Social Sciences

People with roles in life, physical and social sciences. Examples include: professor, chem...

Added on: 04/12/2017

Close Friends of Women with a Birthday in 0-7 days

Close Friends of Women with a Birthday in 0-7 days

Added on: 04/12/2017

Close friends of people with a birthday in 0-7 days

Close friends of people with a birthday in 0-7 days

Added on: 04/12/2017

Architecture and Engineering

Includes Software Engineer, Technician, Electrician, Technical Engineer, Machinist, Archit...

Added on: 03/12/2017

Infrequent attributes Facebook has inferred about you

Attributes that are shared by the smallest number of users on Facebook

Infrequent Interests:



The Interview (2014 film)

News and entertainment

You have this preference because you clicked on an ad related to The Interview (2014 film)...

Nb of facebook users sharing this attribute: 23.79k
Added on: 29/08/2017



Riverbed Technology

Technology

You have this preference because you clicked on an ad related to Riverbed Technology.

Nb of facebook users sharing this attribute: 27.71k
Added on: 03/11/2017

Nicușor Dan

People

You have this preference because you liked a Page related to Nicușor Dan.

Nb of facebook users sharing this attribute: 152.56k
Added on: 24/04/2017



Cloud computing security

Business and industry

You have this preference because you clicked on an ad related to Cloud computing security.

Nb of facebook users sharing this attribute: 166.89k
Added on: 03/11/2017

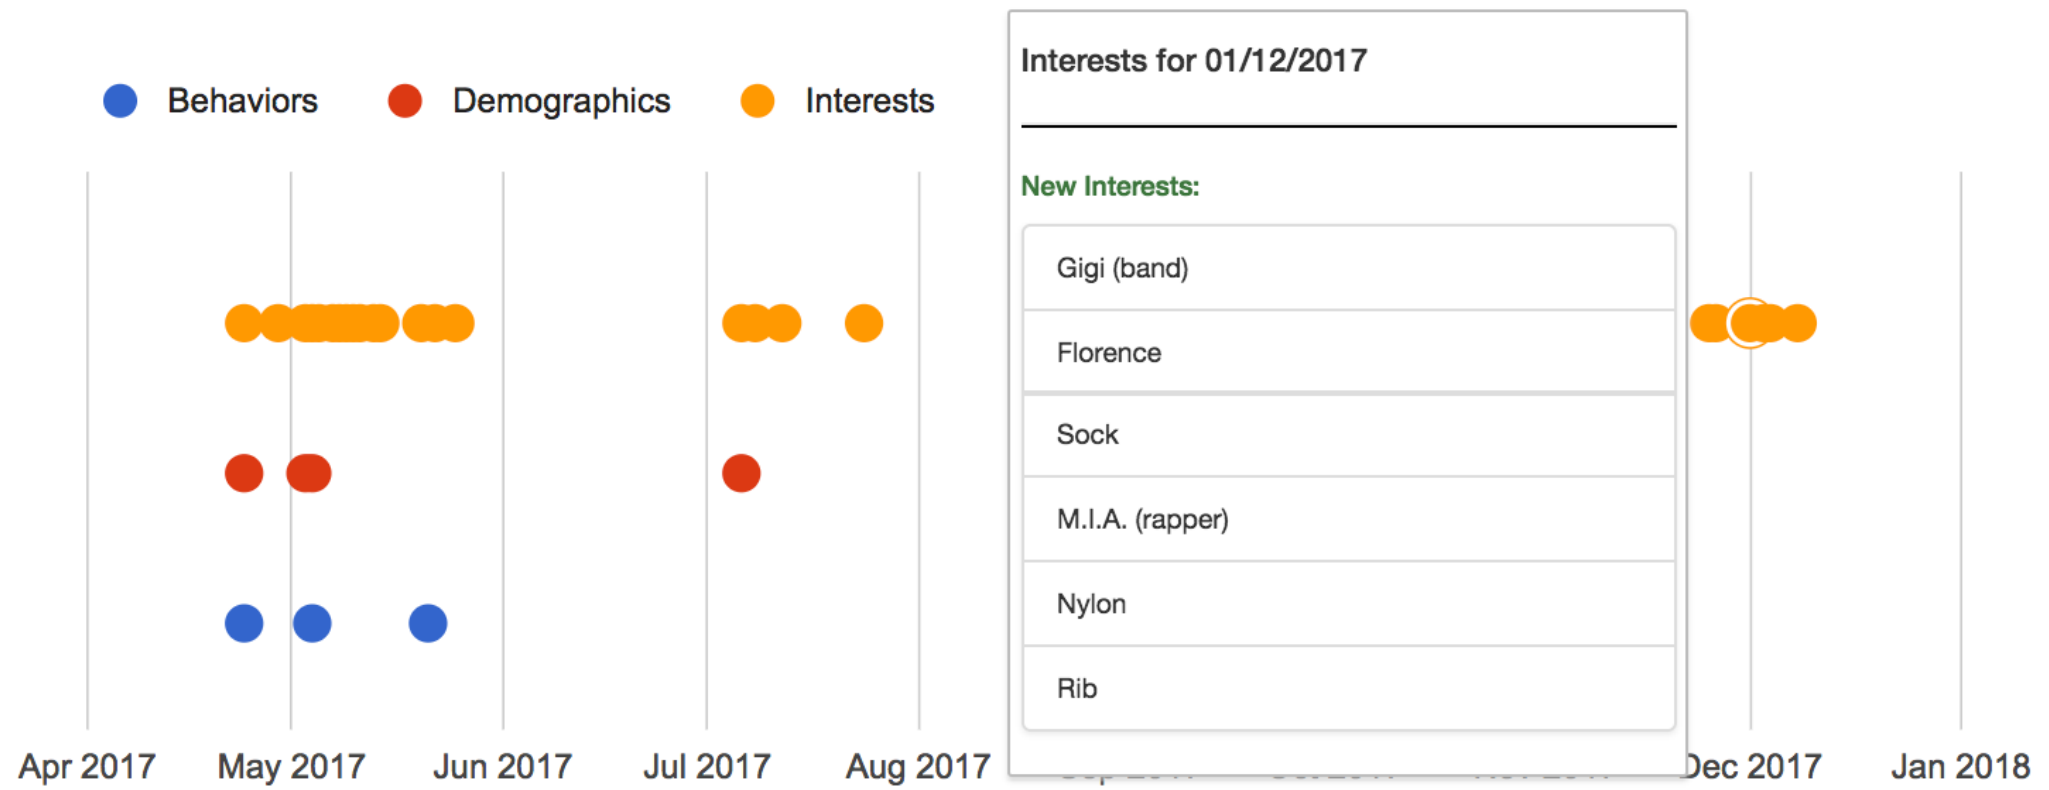
Troll Tennis

News and entertainment

You have this preference because you liked a Page related to Troll Tennis.

Nb of facebook users sharing this attribute: 176.17k
Added on: 24/04/2017

Timeline of when Facebook inferred each attribute:




Top attributes used by advertisers to target you:

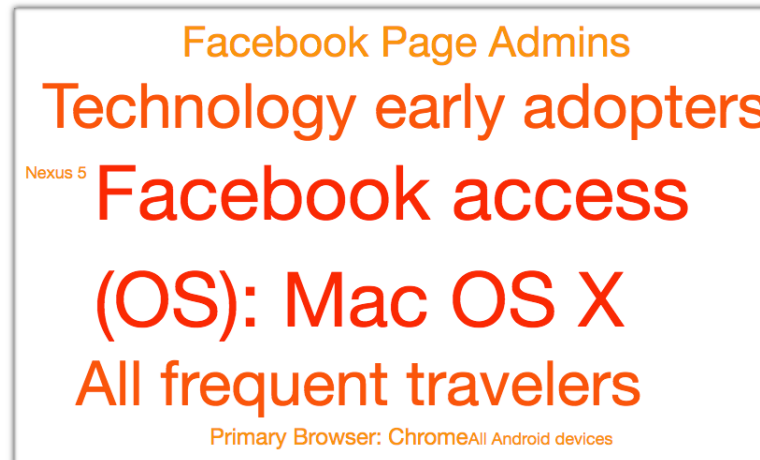
Interests




Demographics



Behaviours



Biographical details



Advertisers view

Latest advertisers targeting you

New advertisers that are targeting you:




Walk Off The Earth
10/12/2017



LE BHV MARAIS
10/12/2017



RED by SFR
09/12/2017



adidas Originals
08/12/2017




Gustave Roussy
07/12/2017

New advertisers that have Personally Identifiable Information about you (advertisers who have your name, phone number, email, address, or whose website or app you have used) that can directly target you:




Panda Curioso
You used the advertiser's website/app.

05/09/2017



Plus.de
You used the advertiser's website/app.

05/09/2017




BuzzFeed Zach
You used the advertiser's website/app.

05/09/2017



Vogue
You used the advertiser's website/app.

03/09/2017




skatedeluxe
You used the advertiser's website/app.

03/09/2017

[Contact us](#) if you didn't give your information to these advertisers

Unpopular advertisers targeting you

Malicious advertisers are rarely popular, here is a list of the advertisers with the lowest number of likes that targeted you




Vacation in Nice
Nb likes: 0



Voyages-Test
Nb likes: 0



Vacation in Saarbrucken
Nb likes: 34



Wordiness
Nb likes: 37



Ginsana Deutschland
Nb likes: 75

Advertisers that use the most unique targeting

Advertisers can target people using infrequent attributes in order to reach a very specific audience (such microtargeting has been used in political ads to influence voters)! These are the advertisers that have targeted you, using the most infrequent attributes.



UNTOLD
Number of users that have the targeting attributes: 342.63k
Targetting attributes: Dub FX



Nature
Number of users that have the targeting attributes: 350.26k
Targetting attributes: Max Planck Society



Chloé
Number of users that have the targeting attributes: 384.08k
Targetting attributes: Hôtel Meurice



Insight Health Apps
Number of users that have the targeting attributes: 487.67k
Targetting attributes: Biofeedback



Paradisía
Number of users that have the targeting attributes: 613.16k
Targetting attributes: Beirut (band)