
Targeted Advertising on Social Media: Transparency, Control, Fairness, Privacy

**Till Speicher, Nina Grigic-Hlaca, Athanasios Andreou,
Giridhari Venkatadri, Oana Goga, Patrick Loiseau,
Krishna P. Gummadi, Alan Mislove**

Max Planck Institute for Software Systems

This talk: Focus on Facebook

- By far, the **largest social media platform**
 - In terms of **number of users**
 - In terms of **data aggregated on users**
 - In terms of **advertisers & ad revenues**
 - In terms of introducing **novel & provocative targeting practices**
 - However, many issues discussed **generalize** to other social media platforms
 - Like LinkedIn, Twitter, YouTube, Instagram etc.,
-

Background: Facebook ad platform

Ads on Facebook



 Spotify

**Spotify Premium
für 0,99 €.
Für alle, die gerne
abtanzen.**

Die Verfügbarkeit des Angebotes ist beschränkt. Es gelten die entsprechenden Nutzungsbedingungen. Danach nur 9,99 €.



€29.39

<http://www.lightinthebox.com/>

Herrenhalbschuhe



Stainless Steel Cooler Stones - 4Pcs

gearbest.com

Free Shipping #CoolerStones #iceball #coolsummer

English (US) · Español (España) · Català ·
Português (Brasil) · Français (France)



Data used for targeting ads

- Facebook gathers lots of data (features) on users

Category	Behaviors														Demographics										Interests							Total					
	<i>Anniversary</i>	<i>Consumer Class.</i>	<i>Digital activities</i>	<i>Expts</i>	<i>Mobile Device User</i>	<i>Multicultural Affinity</i>	<i>Seasonal and Events</i>	<i>Travel</i>	<i>Automotive</i>	<i>Charitable donations</i>		<i>Financial</i>	<i>Job role</i>	<i>Media</i>	<i>Purchase behavior</i>	<i>Residential profiles</i>	<i>Business-to-business</i>	<i>Education level</i>	<i>Generation</i>	<i>Home</i>	<i>Life Events</i>	<i>Parents</i>	<i>Politics (US)</i>	<i>Relationship status</i>	<i>Work</i>	<i>Financial</i>	<i>Business and industry</i>	<i>Entertainment</i>	<i>Family/relationships</i>	<i>Fitness and wellness</i>	<i>Food and drink</i>		<i>Hobbies and activities</i>	<i>Shopping and fashion</i>	<i>Sports and outdoors</i>	<i>Technology</i>	<i>Other</i>
Facebook	1	2	39	74	81	6	2	5									13	3	2	36	9	8	16	26		39	70	8	11	37	60	21	22	21	2	614	
Acxiom								1	5	25	2	35	23	2					19						16												128
Epsilon									4	1			5	2								2															14
Experian													3	1					1																		5
Datalogix							11	151				1	144	29					2		11			1												350	
Total	1	2	39	74	81	6	2	16	152	9	26	3	35	175	5	29	13	3	24	36	20	10	16	27	16	39	70	8	11	37	60	21	22	21	2	1,111	

Closer look at features: Examples

- **Demographical** (gathered by Facebook)
 - **Relationship:**
 - **Interested In:** Men and Women, Men, Unspecified, Women
 - **Status:** Separated, Widowed, Open Relationship, Divorced, In a relationship, Married, Engaged, Unspecified, Single, Complicated Civil Union, Domestic Partnership
 - Each user feature is a **boolean variable**
 - **Demographical** (aggregated from data brokers)
 - **Financial:**
 - **Income:** Geschätztes monatliches Nettoeinkommen 2.000 bis 2.600, 2.600 bis 3.600, 3.600 bis 5.000, über 5.000 Euro
-

Data aggregation across countries

Country	Facebook	Epsilon	DLX	Experian	Acxiom	Total
US	598	14	350	5	128	1105
UK	598	0	19	17	103	737
France	598	0	0	0	21	619
Germany	598	0	0	0	60	658
Australia	598	0	0	34	24	656
Brazil	598	0	0	20	0	618
Japan	598	0	0	0	17	615
South Korea	598	0	0	0	0	598
Canada	598	0	0	0	0	598
India	598	0	0	0	0	598

How Facebook advertisers target users

1) **Traditional targeting**: Advertisers specify a **boolean formula** over the features

- Typically, in a **restricted CNF** form

- $(F_1 \vee F_2 \vee F_3 \dots) \wedge (F'_1 \vee F'_2 \vee F'_3 \dots) \wedge \dots \wedge -F_K \wedge -F'_K$

- Users are targeted, when their **feature values** inferred by Facebook **satisfy the targeting formula**

How Facebook advertisers target users

2) Custom audience targeting: Advertisers upload PII of users they wish to target

Site	Custom audience matching attributes												Minimum Audience	Available targeting attributes								
	First name	Last name	Email	Phone number	ZIP	City	State/Province	Birthday	Gender	Employer	Site user ID	Mobile advertiser ID		Location	Age	Gender	Language	Interests	Behaviors	Mobile device	Keywords	Search terms
Facebook	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	✓	20	✓	✓	✓	✓	✓	✓	✗	✗	✗
Instagram	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓	✓	20	✓	✓	✓	✓	✓	✓	✗	✗	✗
Twitter	✗	✗	✓	✓	✗	✗	✗	✗	✗	✗	✓	✓	500	✓	✗	✓	✓	✓	✓	✓	✓	✗
Google	✗	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	1,000	✓	✗	✗	✓	✗	✗	✓	✓	✓
Pinterest	✗	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	100	✓	✗	✓	✓	✓	✗	✓	✓	✗
LinkedIn	✗	✗	✓	✗	✗	✗	✗	✗	✗	✓	✗	✓	100	✓	✓	✓	✗	✗	✗	✗	✗	✗

- Advertisers love custom audience

This talk: 4 Aspects of Ad Targeting

1. Transparency

2. Control

3. Fairness

4. Privacy

This talk: 4 Aspects of Ad Targeting

1. **Transparency**

- Can I know **what data about me is used to target ads?**

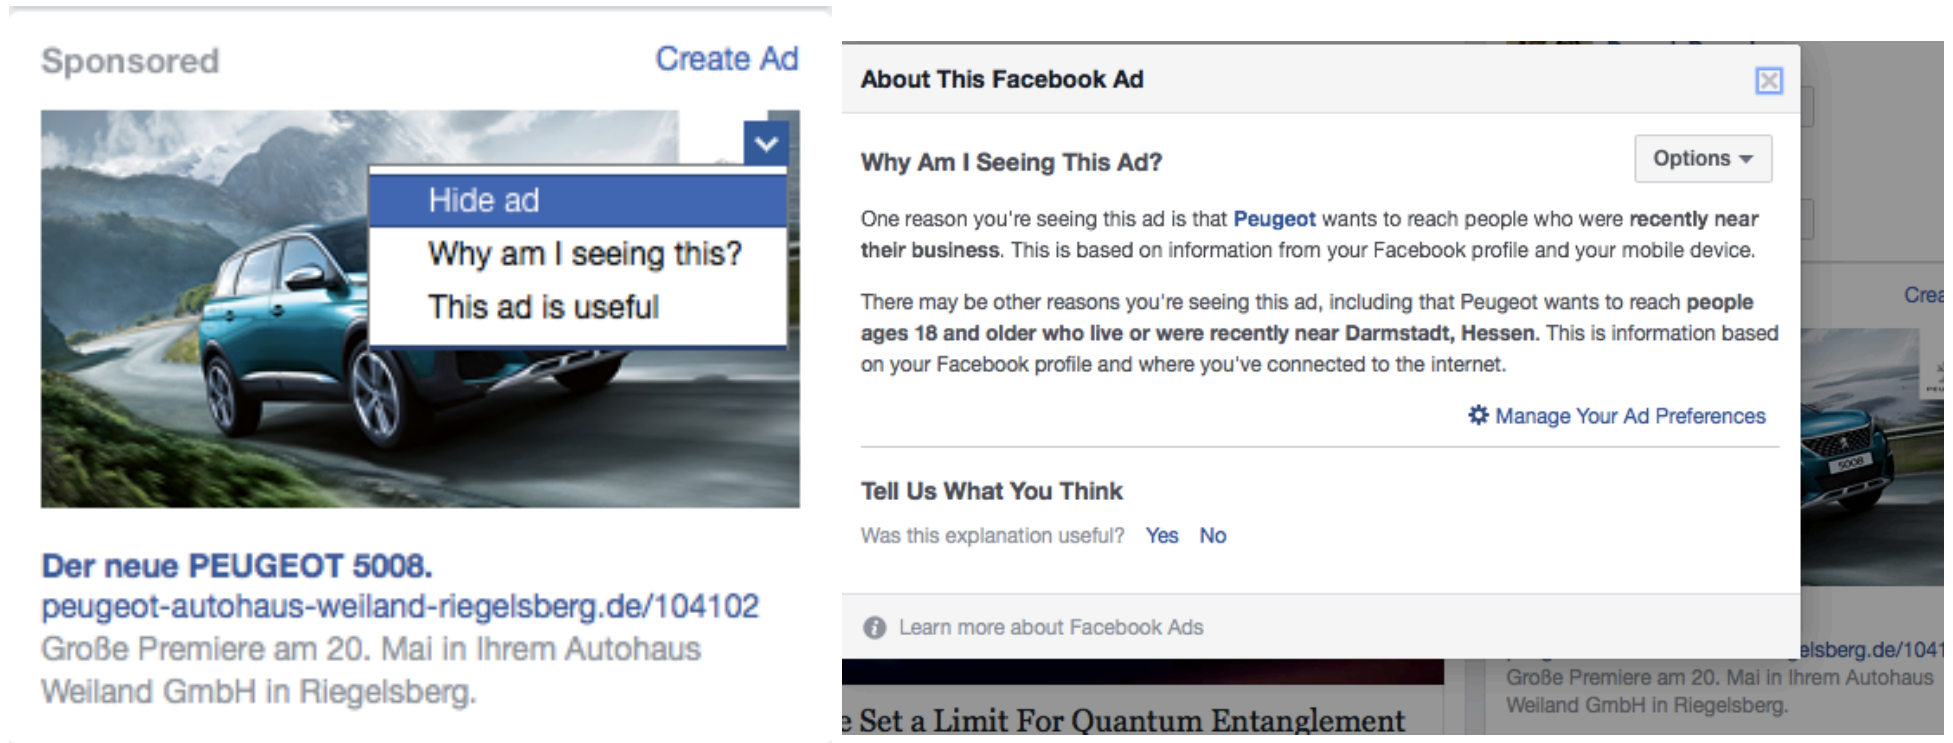
2. **Control**

3. **Fairness**

4. **Privacy**

Transparency of Facebook ad targeting

- Facebook provides transparency via **explanations**



The image shows a Facebook advertisement for a Peugeot 5008. The ad features a teal SUV on a winding road. A dropdown menu is open over the ad, offering options: 'Hide ad', 'Why am I seeing this?', and 'This ad is useful'. To the right, a 'Create Ad' button is visible. Below the ad, the text reads: 'Der neue PEUGEOT 5008. peugeot-autohaus-weiland-riegelsberg.de/104102 Große Premiere am 20. Mai in Ihrem Autohaus Weiland GmbH in Riegelsberg.'

The 'About This Facebook Ad' overlay provides the following information:

- Why Am I Seeing This Ad?** (Options ▾)
- One reason you're seeing this ad is that **Peugeot** wants to reach people who were **recently near their business**. This is based on information from your Facebook profile and your mobile device.
- There may be other reasons you're seeing this ad, including that Peugeot wants to reach **people ages 18 and older who live or were recently near Darmstadt, Hessen**. This is information based on your Facebook profile and where you've connected to the internet.
- [Manage Your Ad Preferences](#)
- Tell Us What You Think**
- Was this explanation useful? Yes No
- [Learn more about Facebook Ads](#)

- Largely **voluntary**
 - At times to **satisfy legal requirements**

How complete are the explanations?

- ❑ For ads targeting **customer PII**:
 - ❑ *"One reason you're seeing this ad is that [Booking.com](#) wants to reach people who have visited their website or used one of their apps. This is based on [customer information](#) provided by [Booking.com](#)."*
 - ❑ Doesn't state **what PII Booking.com used!**
 - ❑ E.g., Did booking.com use email? Or phone number? Or name/address?
-

How complete are the explanations?

- For ads targeting **data broker info**:
 - *"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."*
 - Doesn't state **what Acxiom provided data** was used!
 - E.g., is it based on financial data? Or purchasing habits?
-

How complete are the explanations?

- For ads targeting **FB user data**:
 - Beyond location, gender, age: picks **exactly one** of the several features used in targeting formula
 - *"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."*
-

Validation of incompleteness

- ❑ Ran a **controlled ad** targeting ourselves using a **custom list** and selecting **millennials & 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.."*
 - ❑ *"There may be other reasons why you're seeing this advert, including that Acer 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."*
- ❑ Unclear **why millennials** was chosen **over expats!**

Do explanations need to be complete?

- Should they specify **all features in ad targeting?**
 - Arguments **for**:
 - **Avoids misleading / fake** 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 features for explanations

- Example explanation:

"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 features **the most important**?
 - Is Shopping and fashion the most important of all the user' features that Facebook and the advertised used to target the user?
 - **Not clear what criterion Facebook is using!**
-

Open challenges

- How to pick **a few (K) features** for explanations?
 - How to determine the **importance of a user feature**?
 - Does it reveal **privacy sensitive information**?
 - Is it a **rare (or low prevalence) feature in population**?
 - Does it exert **the most influence**?
 - On sizing the target audience?
 - As measured by “Quantifying Input Influence” framework?
-

This talk: 4 Aspects of Ad Targeting

1. Transparency

- ❑ Can I know **what data about me is used to target ads?**
- ❑ **NO! Only some (not all)** data used is revealed

2. Control

3. Fairness

4. Privacy

This talk: 4 Aspects of Ad Targeting

1. Transparency

- ❑ Can I know **what data about me is used to target ads?**
- ❑ **NO! Only some (not all)** data used is revealed

2. Control

- ❑ Do explanations offer consumers **satisfactory control?**

3. Fairness

4. Privacy

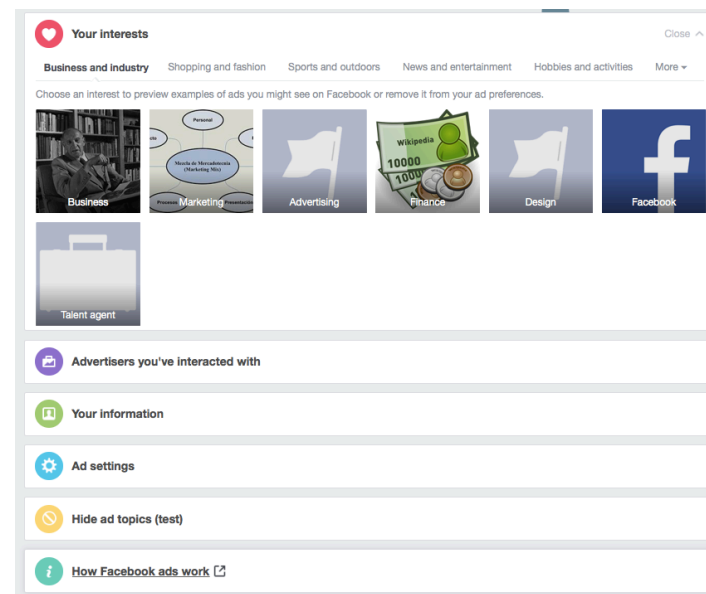
Consumer control over targeted ads

- Control **Inputs**: Determine **data used for targeting** ads
 - Help consumers protect their privacy

 - Control **Outputs**: Determine **ads they receive**
 - Help consumers achieve specific decision outcomes
-

Consumer controls in Facebook

- Facebook claims that **ad preferences page** allows **control** over what ads users see



- In reality, it offers **controls over inputs, not outputs!**
-

Do explanations offer output control?

- Given the **incompleteness** of explained features:

"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 features **necessary**?
 - If the user asked Facebook to **remove their data on shopping and fashion**, would they **stop receiving the ad**?
 - **No guarantees!**
-

Do explanations offer output control?

- Given the **incompleteness** of explained features:

"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 features **sufficient**?
 - If the user asked Facebook to **change their data on features other than shopping and fashion**, would they continue to receive the ad?
 - **No guarantees!**
-

Open challenges

- Given the **incompleteness** of explained features
 - How to construct **necessary/sufficient** explanations?
 - Technically **challenging** for arbitrary targeting formulae!
 - Finding **min. set of necessary or sufficient features**
 - Often, maps to solving **NP-Complete SAT** problem
 - But, can leverage recent **advances on SAT solvers**!
-

This talk: 4 Aspects of Ad Targeting

1. Transparency

- ❑ Can I know **what data about me is used to target ads?**
- ❑ **NO! Only some (not all)** data used is revealed

2. Control

- ❑ Do explanations offer consumers **satisfactory control?**
- ❑ **NO!** Constructing satisfactory explanations is **non-trivial**

3. Fairness

4. Privacy

This talk: 4 Aspects of Ad Targeting

1. Transparency

- ❑ Can I know **what data about me is used to target ads?**
- ❑ **NO!** Only some (not all) data used is revealed

2. Control

- ❑ Do explanations offer consumers **satisfactory control?**
- ❑ **NO!** Constructing satisfactory explanations is **non-trivial**

3. Fairness

- ❑ Can we **detect & prevent** discriminatory ad targeting?

4. Privacy

Can ad targeting be discriminatory?



[Online Ads for High-Paying Jobs Are Targeting Men More Than ...](#)

Adweek - 7 Jul 2015

"We found small instances where there was **discrimination** and **gender-based discrimination** in job ads," said ... The issue of bias and **discrimination** in ad targeting isn't new, ... The **Carnegie Mellon researchers** also were alarmed by ads that ... (Google's rules forbid serving ads based on health information.)

[Carnegie Mellon Study Finds Gender Discrimination In Ads Shown ...](#)

Marketing Land - 8 Jul 2015

[When Algorithms Discriminate](#)

New York Times - 9 Jul 2015

[Probing the Dark Side of Google's Ad-Targeting System](#)

Highly Cited - MIT Technology Review - 6 Jul 2015

[Google's Ad System Has Become Too Big to Control](#)

In-Depth - Wired - 9 Jul 2015

[Google's algorithm shows prestigious job ads to men, but not to ...](#)

Highly Cited - Washington Post - 6 Jul 2015



Marketing Land



New York Tim...



The Guardian



Washington P...



Chicago Tribu...



TechCrunch

[View all](#)

Can ad targeting be discriminatory?



[Facebook To Ban 'Ethnic Affinity' Targeting For Housing ...](#)

Forbes - 11 Nov 2016

After weeks of push back from U.S. lawmakers, media and civil rights leaders, Facebook FB +0.18% on Friday announced it will stop allowing ...

[Facebook disables 'ethnic affinity' ads for housing, jobs](#)

Engadget - 11 Nov 2016

[Facebook to stop ads that target, exclude races](#)

Highly Cited - USA TODAY - 11 Nov 2016

[Facebook to Remove 'Ethnic Affinity' Targeting From Certain Ad ...](#)

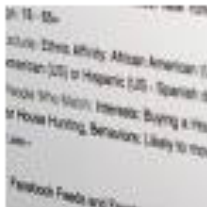
International - AgencySpy - 11 Nov 2016

[Facebook has discriminated against you, and it's not going to stop](#)

In-Depth - Mashable - 12 Nov 2016

[Facebook Bans Targeting Based on Race and Ethnicity for Housing ...](#)

Blog - Slate Magazine (blog) - 11 Nov 2016



Engadget



TechCrunch



Deadline



Daily Mail



SlashGear



USA TODAY

[View all](#)

Discrimination via correlated features

- FB's early defense: **Ethnic affinity is not ethnicity**
- Used **voter records from NC** to check correlations
 - Voter records have race information
 - Created separate customer lists for different races
 - Checked correlations between their race & ethnic affinity

Attribute	Voter Records		Facebook Users					
	Number	Percent	Uploaded	Matched	Reachable	Reachable %	Corresponding	Corresponding %
White	5,303,383	70.1%	10,000	8,000	6,800	85.0%	5,700	83.8%
Black	1,694,220	22.4%	10,000	7,800	6,300	80.8%	5,200	82.5%
Asian	79,250	1.0%	10,000	7,700	6,600	85.7%	1,900	28.8%
Hispanic	163,236	2.2%	10,000	7,000	5,900	84.3%	3,000	50.8%

Does banning “ethnic affinity” help?

- What about **pre-filtered custom lists**:
 - **using offline info** like voter records?
- What of other **correlated features**?

no	feature name	selectivity	Blacks percentage	rest percentage	ratio
1	Demographic > Ethnic Affinity > African American (US)	17.0%	77.0%	10.9%	7.06
2	Demographic > Politics (US) > US Politics (Very Liberal)	11.8%	49.8%	7.7%	6.44
3	Interests > Entertainment > Music > Gospel music	14.4%	48.3%	14.6%	3.32
4	Interests > Shopping and fashion > Beauty > Hair products	12.2%	40.8%	12.9%	3.15

no	feature name	selectivity	Blacks percentage	rest percentage	ratio
1	Demographic > Politics (US) > US Politics (Very Conservative)	14.4%	4.8%	26.5%	0.18
2	Demographic > Politics (US) > US Politics (Conservative)	16.6%	6.5%	29.7%	0.22
3	Interests > Sports and outdoors > Outdoor recreation > Hiking	11.0%	8.0%	21.6%	0.37
4	Interests > Sports and outdoors > Outdoor recreation > Camping	11.4%	11.5%	22.8%	0.50

Open challenges

- ❑ How to **detect discriminatory targeting** in ads?
 - ❑ Particularly, with customer lists?
 - ❑ How to **avoid discriminatory targeting** in ads?
 - ❑ Detecting & avoiding algo. discrimination is a **hot topic**
 - ❑ But, even here ads pose **unique challenges**
 - ❑ **Fair targeting** might result in **unfair ad impressions!**
 - ❑ Targeting 100 men & 100 women might result in unequal impressions, when **costs of their impressions are different!**
-

This talk: 4 Aspects of Ad Targeting

1. Transparency

- ❑ Can I know **what data about me is used to target ads?**
- ❑ **NO!** Only some (not all) data used is revealed

2. Control

- ❑ Do explanations offer consumers **satisfactory control?**
- ❑ **NO!** Constructing satisfactory explanations is **non-trivial**

3. Fairness

- ❑ Can we **detect & prevent** discriminatory targeting?
- ❑ **Unclear** – need **new measures & methods** for fairness

4. Privacy

This talk: 4 Aspects of Ad Targeting

1. Transparency

- ❑ Can I know **what data about me is used to target ads?**
- ❑ **NO!** Only some (not all) data used is revealed

2. Control

- ❑ Do explanations offer consumers **satisfactory control?**
- ❑ **NO!** Constructing satisfactory explanations is **non-trivial**

3. Fairness

- ❑ Can we **detect & prevent** discriminatory targeting?
- ❑ **Unclear** – need **new measures & methods** for fairness

4. Privacy

- ❑ Are there any **privacy risks** with targeted ads?
-

Potential audience reach estimate

- A feature of Facebook's advertiser interface

Audience
Define who you want to see your ads. [Learn more.](#)

[Create New](#) Use a Saved Audience ▾

Custom Audiences ⓘ
Customer List
user-gmails
Add Custom Audiences or Lookalike Audiences
Exclude | Create New ▾

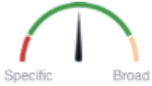
Locations ⓘ
Everyone in this location ▾
United States
📍 **United States**
📍 Include ▾ | Add locations
Add Bulk Locations...

Age ⓘ
18 ▾ - 65+ ▾

Gender ⓘ
All Men Women

Languages ⓘ
Enter a language...

Detailed Targeting ⓘ
INCLUDE people who match at least ONE of the following ⓘ
Add demographics, interests or behaviors | Suggestions | Browse
Exclude People



Audience Definition

Your audience is defined.



Audience Details:

- Custom Audience:
 - user-gmails
- Location:
 - United States
- Age:
 - 18 - 65+
- Placements:
 - Facebook Feeds and Instagram Feed

Potential Reach: 200,000 people

Estimated Daily Reach

 **1,800 - 4,800 people on Facebook**
0  of 160,000 ⓘ

 **920 - 2,400 people on Instagram**
0  of 66,000 ⓘ

This is only an estimate. Numbers shown are based on the average performance of ads targeted to your selected audience.

Privacy risks from audience estimates

- Assume **exact estimates** of audience size
 - Then, **given a user's PII**
 - Like phone-num. or email-id. or name-address
 - Any advertiser can **check if the user is on Facebook**
 - And **retrieve all the user's info** Facebook inferred
 - Including the financial info provided by data brokers!
-

Precision of audience reach estimates

- **Reverse-engineered** how the estimates work
 - **No estimates** given when the **audience reach** < 20
 - **Estimates are rounded**
 - Audience reach < 1000, rounded down to closest 10
 - Audience reach < 10000, rounded down to closest 100
 - Audience reach < 100000, rounded down to closest 1000
 -
 - **Estimates are transient, but stable over short-term**
-

Privacy risks from audience estimates

- Given any customer list S with and a user U 's PII
 - Like phone-num. or email-id. or name-address
 - Create a new customer list with $S + U$
 - Is audience reach for $S + U$ is more than S ?
 - If it does, user U is on FB
 - One can similarly retrieve all the info FB has on the user
 - If not, either U is not on FB OR it's a rounding error
-

The probability of rounding errors

- For a list S with audience reach > 20 & < 1000
 - Rounding error probability is 0.9
 - For K -lists with audience reach > 20 & < 1000
 - Chance of every try suffering a rounding error is 0.9^K
 - Chance of at-least one try not being rounded is $1 - 0.9^K$
 - For $K = 100$, this is chance is 99.999974%
 - So **by creating 100 lists** with reach > 20 & < 1000
 - One can **w.h.p. retrieve all data FB has on any user**
-

How to create such customer lists

- Use **public voter records** in the US!
 - Randomly sample names/addresses from records
 - Till you get a customer list of size > 20 & < 1000 !
 - Repeat the process 100 times!

 - **Validation**: Used it to retrieve all data FB has on us!
 - **User transparency tool**?
 - But, could be used to **retrieve data on others** as well!
-

Open challenges

- Audience estimates are very **useful for advertisers**
 - How to preserve the **estimates without data leaks**
 - **Differential privacy? Other noisy estimates?**
 - But, how can one enforce privacy budgets?
 - Other privacy risks with custom lists?
 - **Creepy ads?**
 - E.g., targeting sex-offenders using public criminal records
 - More potent attacks?
-

Summary: Ad Targeting on Social Media

1. Transparency

- ❑ Limited! Only some (not all) data used for is revealed

2. Control

- ❑ Limited! Providing satisfactory explanations is non-trivial

3. Fairness

- ❑ Unclear! Need new measures & methods for fairness

4. Privacy:

- ❑ Non-existent! Avoiding data leaks via ads is non-trivial!

- ❑ Lots of open challenges!
-