

# Echoes of Privacy: Uncovering the Profiling Practices of Voice Assistants

*Tina Khezresmaeilzadeh<sup>1</sup>, **Elaine Zhu**<sup>2</sup>, Kiersten Grieco<sup>2</sup>,  
Daniel J. Dubois<sup>2</sup>, Konstantinos Psounis<sup>1</sup>, David Choffnes<sup>2</sup>*

*<sup>1</sup>University of Southern California    <sup>2</sup>Northeastern University*



**USC** University of  
Southern California



Northeastern  
University

# The Problem

142M+

Voice Assistant Users in USA

3

Major Platforms

- ▶ Voice assistants are applications that use voice technology to carry out tasks to help users
- ▶ Voice interactions contain potentially sensitive and revealing data
- ▶ Privacy policies mention profiling but **lack specifics**
- ▶ **Key Question:** How do voice assistants actually profile users?

# What constitutes profiling?

Collecting data on individuals or groups based on specific traits or behaviors to understand them.

- **Recommendations and personalization.** We use your personal information to recommend features, products, and services that might be of interest to you, identify your preferences, and personalize your experience with Amazon Services.

Collecting data on individuals or groups based on specific traits or behaviors to understand them.

- **Recommendations and personalization.** We use your personal information to recommend features, products, and services that might be of interest to you, identify your preferences, and personalize your experience with Amazon Services.

## Collecting data on individuals or groups

If ad personalization is turned on, Google will use your information to make your ads more useful for you. For example, a website that sells mountain bikes might use Google's ad services. After you visit that site, you could see an ad for mountain bikes on a different site that shows ads served by Google.

# Research Questions

# Research Questions

## RQ2

Does profiling actually happen and to what extent?

# Research Questions

## RQ1

Do voice assistants have  
prepopulated labels  
before any interaction?

## RQ2

Does profiling actually  
happen and to what  
extent?



# Research Questions

## RQ1

Do voice assistants have **prepopulated labels** before any interaction?

## RQ2

Does **profiling actually happen** and to what extent?

## RQ3

Does **interaction modality** (voice vs. web) affect profiling?

# Research Questions

## RQ1

Do voice assistants have **prepopulated labels** before any interaction?

## RQ2

Does **profiling actually happen** and to what extent?

## RQ3

Does **interaction modality** (voice vs. web) affect profiling?

## RQ4

Are **opt-out mechanisms** effective?

# Challenges and Experimental Design

## **Fresh Accounts**

Unique phone  
numbers, factory  
reset devices

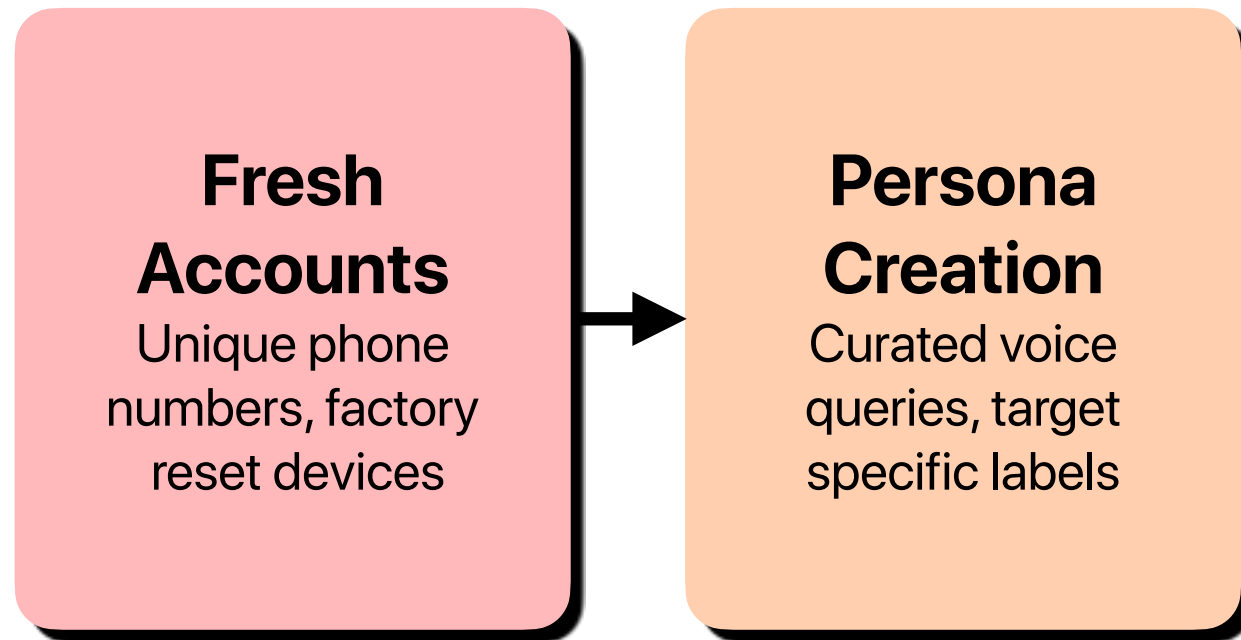
## **Controlled Environment**

Same IP address,  
automated playback,  
supervised execution

## **Persona-Based Approach**

10 accounts per persona,  
query generation guided by  
ChatGPT-4 with manual  
supervision

# Challenges and Experimental Design



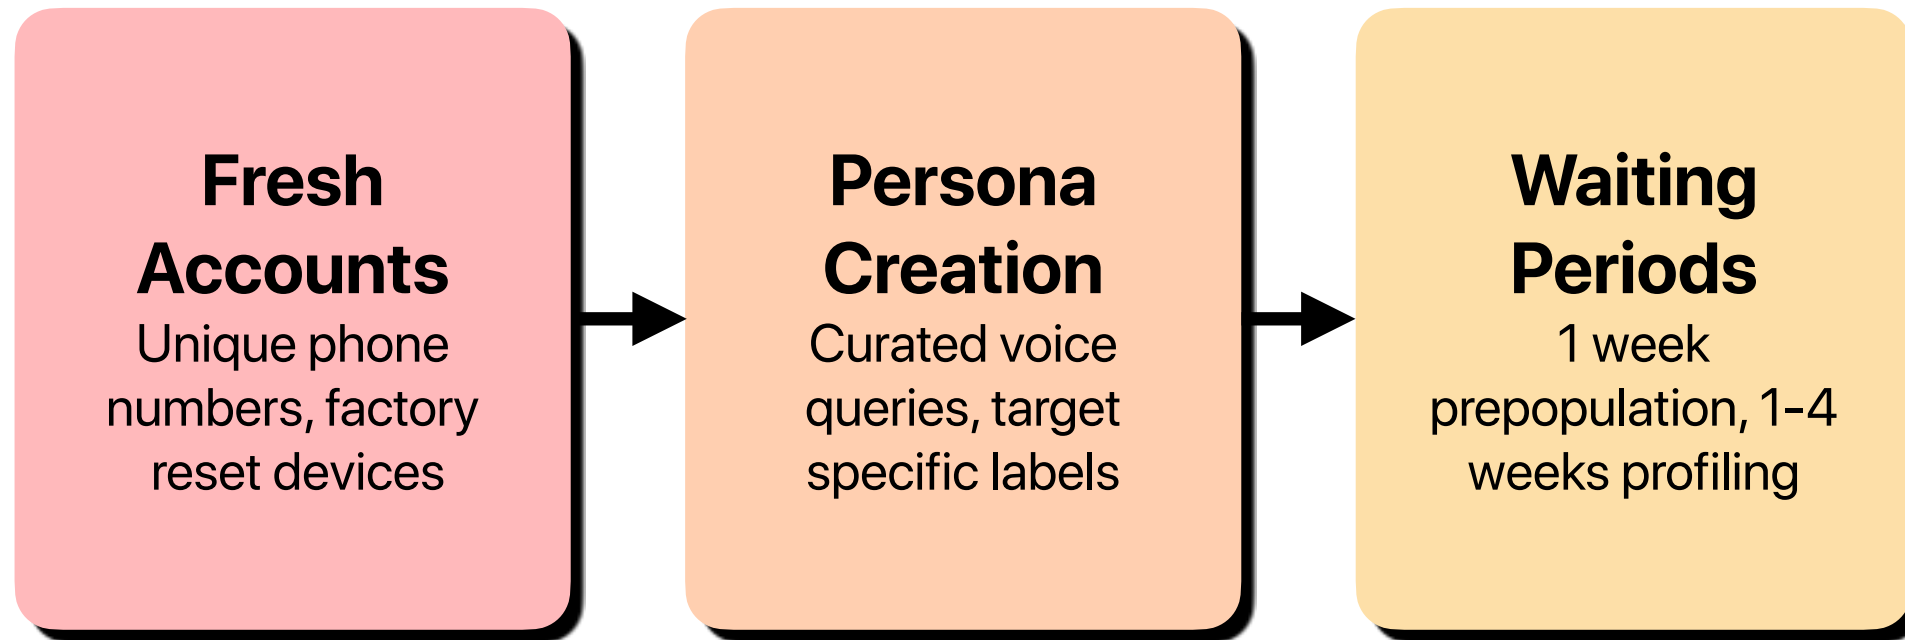
## Controlled Environment

Same IP address,  
automated playback,  
supervised execution

## Persona-Based Approach

10 accounts per persona,  
query generation guided by  
ChatGPT-4 with manual  
supervision

# Challenges and Experimental Design



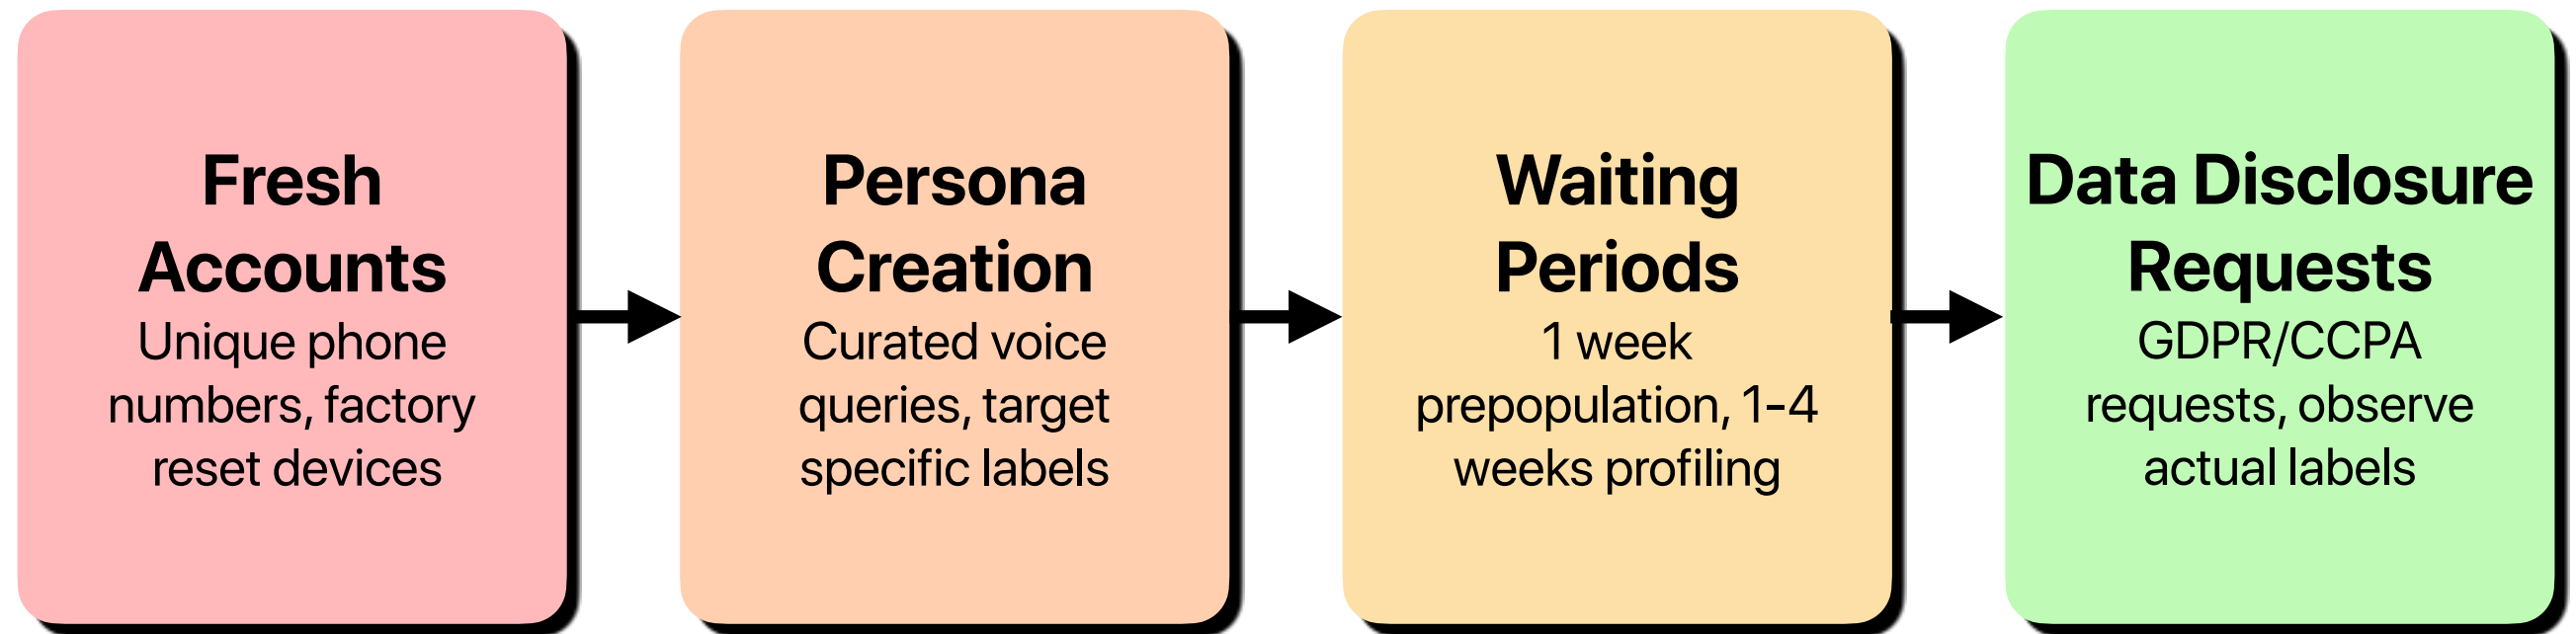
## Controlled Environment

Same IP address,  
automated playback,  
supervised execution

## Persona-Based Approach

10 accounts per persona,  
query generation guided by  
ChatGPT-4 with manual  
supervision

# Challenges and Experimental Design



## Controlled Environment

Same IP address,  
automated playback,  
supervised execution

## Persona-Based Approach

10 accounts per persona,  
query generation guided by  
ChatGPT-4 with manual  
supervision

# Summary of Experiments

1,171

Experiments

24,530

Voice Queries

20

Months Duration

# Summary of Experiments

1,171


Experiments

24,530

Voice Queries

20

Months Duration



19 Personas

*Demographic Labels*



# Summary of Experiments

1,171


Experiments

24,530

Voice Queries

20

Months Duration



19 Personas  
*Demographic Labels*




17 Personas  
*Interest Labels*

# Summary of Experiments

1,171  
Experiments

24,530  
Voice Queries

20  
Months Duration



19 Personas  
*Demographic Labels*




17 Personas  
*Interest Labels*




19 Personas  
*No Direct Profiling*

# **Key Finding 1: Prepopulated Labels (RQ1)**

# Key Finding 1: Prepopulated Labels (RQ1)



# Key Finding 1: Prepopulated Labels (RQ1)




195/200 Accounts:  
"Homeowners" + "Not  
Parents"  
194/200 Accounts:  
"Moderately High  
Income"



No prepopulated  
labels

# Key Finding 1: Prepopulated Labels (RQ1)



195/200 Accounts:  
"Homeowners" + "Not  
Parents"  
194/200 Accounts:  
"Moderately High  
Income"




No prepopulated  
labels



No prepopulated  
labels

# Key Finding 1: Prepopulated Labels (RQ1)



195/200 Accounts:  
"Homeowners" + "Not  
Parents"  
194/200 Accounts:  
"Moderately High  
Income"



No prepopulated  
labels



No prepopulated  
labels

- ▶ **Privacy Implication:** profiling without interaction
- ▶ Labels could be **inaccurate** and affect ad targeting
- ▶ Mechanisms unclear

# **Key Finding 2: Profiling Characterization (RQ2)**



# Key Finding 2: Profiling Characterization (RQ2)

Platform	Profiling Type (Method used to assign labels)	Accuracy (Percentage of correctly assigned target labels)
Google	Probabilistic	10-80% (varies)
Amazon	Deterministic	100% (commands only)
Apple	None detected	N/A

# Key Finding 2: Profiling Characterization (RQ2)

Platform	Profiling Type (Method used to assign labels)	Accuracy (Percentage of correctly assigned target labels)	Time to Profile (Days required for labels to appear after queries)
Google	Probabilistic	10-80% (varies)	18.0 ± 4.1 days
Amazon	Deterministic	100% (commands only)	7.7 ± 1.3 days
Apple	None detected	N/A	N/A

# Key Finding 2: Profiling Characterization (RQ2)

Platform	Profiling Type (Method used to assign labels)	Accuracy (Percentage of correctly assigned target labels)	Time to Profile (Days required for labels to appear after queries)	Consistency (Reliability of getting same results across experiments)
Google	Probabilistic	10-80% (varies)	18.0 ± 4.1 days	Inconsistent
Amazon	Deterministic	100% (commands only)	7.7 ± 1.3 days	Consistent
Apple	None detected	N/A	N/A	N/A

# Key Finding 2: Profiling Characterization (RQ2)

Platform	Profiling Type (Method used to assign labels)	Accuracy (Percentage of correctly assigned target labels)	Time to Profile (Days required for labels to appear after queries)	Consistency (Reliability of getting same results across experiments)
Google	Probabilistic	10-80% (varies)	18.0 ± 4.1 days	Inconsistent
Amazon	Deterministic	100% (commands only)	7.7 ± 1.3 days	Consistent
Apple	None detected	N/A	N/A	N/A

- ▶ Google accuracy decreased over time: 80% → 20% for same labels
- ▶ Amazon only profiles shopping commands, not general queries
- ▶ Profiling timing creates transparency gap
- ▶ **Privacy Implication:** Opaque profiling with mislabeling risk

# **Key Finding 3: Modality Effects (RQ3)**

# Key Finding 3: Modality Effects (RQ3)




**Google Web:**  
62.86% accuracy  
2.2 days average

# Key Finding 3: Modality Effects (RQ3)



**Google Web:**  
62.86% accuracy  
2.2 days average




**Google Voice:**  
48% accuracy  
18.0 days average

# Key Finding 3: Modality Effects (RQ3)



**Google Web:**  
62.86% accuracy  
2.2 days average



**Google Voice:**  
48% accuracy  
18.0 days average




**Amazon (both):**  
100% accuracy  
7-8 days average



# Key Finding 3: Modality Effects (RQ3)



**Google Web:**  
62.86% accuracy  
2.2 days average



**Google Voice:**  
48% accuracy  
18.0 days average



**Amazon (both):**  
100% accuracy  
7-8 days average

- ▶ **Google:** Web modality profiles faster and more accurately
- ▶ **Amazon:** Similar performance across modalities
- ▶ **Privacy implication:** Voice users face higher mislabeling risk when using Google
- ▶ Suggests different profiling algorithms for different modalities

# Key Finding 4: Mitigation Effectiveness (RQ4)






All platforms offer opt-out mechanisms, but only Google provides granular control.

**Privacy Implication:** Other than Google, limited user control for over incorrect labels.

# Key Finding 4: Mitigation Effectiveness (RQ4)

Categories used to show you ads

You may see ads meant for people in these categories, which are based on your Google activity. You can get ads for a different category, or turn off anything you don't want used.

 On	 On	 On	 On	 On
Relationships In a relationship	Household Income Moderately high i...	Education Bachelor's degree	Industry Not enough info	Employment Large

Google

**Privacy Implication:** Other than Google, limited user control for over incorrect labels.

# Key Finding 4: Mitigation Effectiveness (RQ4)

## Categories used to show you ads

You may see ads meant for people in these categories, which are based on your Google activity. You can get ads for a different category, or turn off anything you don't want used.



On

Relationships  
In a relationship



On

Household Income  
Moderately high i...



On

Education  
Bachelor's degree



On

Industry  
Not enough info



Employment  
Large

Google

## Privacy Implication control for

### Turn off Web & App Activity for ads?

Web & App Activity won't be used to personalize ads

This means that your activity, including things you've searched for and sites you've visited, won't be used to influence which ads you see. Your Web & App Activity may still be saved.

Ads may seem less relevant

When your activity isn't used to personalize your ads, fewer of the ads you see may be about products and brands that interest you.

This setting affects personalized ads on:



Search



YouTube



Discover

Cancel

Turn off

# Privacy Implications

## Transparency Issues

- Timing delays obscure activity
- Undisclosed profiling methods
- Prepopulated labels without knowledge
- Probabilistic profiling → mislabeling

## User Control Limitations

- Limited granular control (Amazon)
- Modality affects outcomes

**Gap between policies and actual practices**

# Recommendations

- ▶ Leverage available opt-out mechanisms
- ▶ Increase transparency about profiling timelines
- ▶ Provide confidence levels for labels
- ▶ Strengthen disclosure requirements

\*  = for users     = for platforms     = for regulators

# Conclusion and Future Work

## Key Contributions

- First comprehensive study of voice assistant profiling
- Platform-specific differences in profiling practices
- Modality effects on profiling accuracy and timing

## Future Directions

- Multi-language profiling analysis
- Long-term profiling evolution studies
- Third-party skill integration effects
- Voice characteristics profiling detection



Artifacts can be found at:  
<https://tinyurl.com/voiceassistants2025>  
or the QR code on this slide



Questions? Email me at  
[elainezhu@ccs.neu.edu](mailto:elainezhu@ccs.neu.edu)!