



UNDERSTANDING AND CAPTURING PEOPLE'S MOBILE APP PRIVACY PREFERENCES

Jiali Lin
jialiul@cs.cmu.edu
School of Computer Science
Carnegie Mellon University

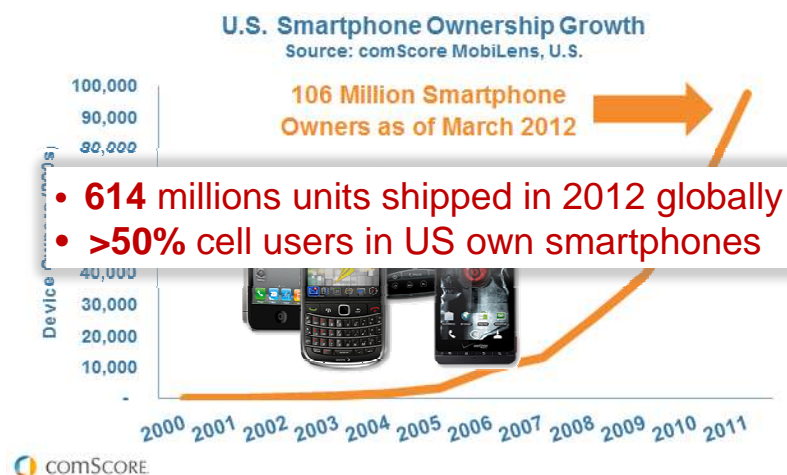
Committee
Jason I. Hong (co-chair)
Norman Sadeh (co-chair)
Mahadev Satyanarayanan
Sunny Consolvo

Thesis Proposal, Aug 17, 2012

SCHOOL OF COMPUTER SCIENCE
CARNEGIE MELLON UNIVERSITY



Smartphone Ownership Rockets



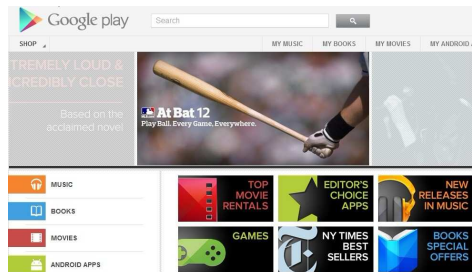
<http://blog.comscore.com/2012/05/comscore-introduces-mobile-matrix-20-the-next-gene.html>

2/50

Mobile App Stores Thrive



- By June 2012, Apple offered **650,000** apps with **30 billion** downloads (WWDC, 2012)



- Google Play Store offered **600,000** apps with more than **20 billion** downloads (Google I/O 2012)

3/50

Mobile Apps Also Bring Privacy Risks

- Lots of mobile apps use private information in surprising ways



Pandora gathers location, gender, year of birth, etc.



Path uploads entire contact list without user full consent.



Brightest Flashlight requires full Internet access, location, etc.




Bible accesses location.

- A significant portion of private information goes to 3rd parties through APIs



4/50

Existing Privacy Interfaces and Settings Fall Short



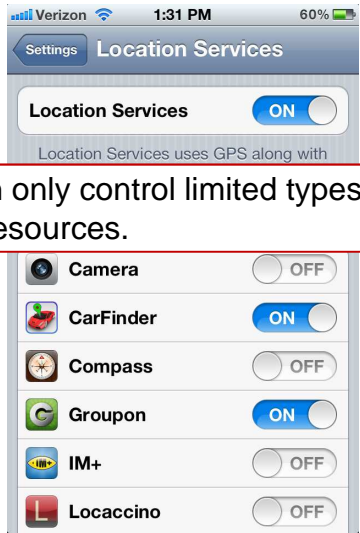
PERMISSIONS

- Need to review the permissions before installing the app.
- Cannot tell where private info goes.

Hardware controls

People do not read permissions or cannot understand them. (Felt'12, Kelley'12),

Network communication



Can only control limited types of resources.

5/50

Where My Research Fits

	Context Sharing Services	Mobile Apps in General
Technological Solution	Computational privacy Trusted Middleware	App analysis Security extensions
User Research	Location privacy policy learning Place Naming Identify influencing factors US-China Study Essential privacy controls Purpose- vs. social-driven	<div style="border: 2px solid red; border-radius: 15px; padding: 10px;"> How to understand app behaviors from users' perspectives How to model users' mobile app privacy preferences </div> Expectation & Purpose Permission usability study Survey of mobile privacy concerns
	 My past work	 My proposed work

6/50

Thesis Statement

Thesis Statement:

*By using **crowdsourcing** and **user-oriented machine learning** techniques, we can build accurate and understandable models of mobile apps and users' privacy preferences to inform the design of mobile privacy interfaces and settings, and to help developers build more privacy preserving apps.*

7/50

More Specifically...

- Bridge the gap between technological solutions and user research
- Experiment new ways, i.e., **crowdsourcing**, to capture users' mobile app privacy preferences
- Model users' preferences through **user-oriented ML techniques** to generalize our findings

8/50

Research Contributions

- A new way of looking at mobile privacy, i.e. “Privacy as expectation”
- A valuable dataset containing both app behavioral attributes and users privacy ratings
- Clusters of mobile apps elicit distinct privacy concerns
- A set of **privacy personas** that can simplify privacy settings
 - i.e., common privacy policies shared by a group of users

9/50

Expected Benefits of this Research

- Provide **mobile app markets** models to evaluate apps from a privacy perspective
- Inform the design of usable and efficient privacy interfaces and settings of existing **mobile OS**
- Help **developers** to understand the privacy implications and user acceptance of their apps






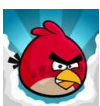
10/50

Outline

- Background
- **Related Work in Location Sharing**
- Related Work in Mobile App Privacy
- Proposed Work: Investigating Users' Mobile App Privacy Preferences
- Schedule and Summary

11/50

Different Context of Use

	Location-sharing Apps	Mobile Apps in General
Entity Accessing Sensitive Resources	Other end-users	Primary: Service providers, Secondary: advertising agent, etc.
Examples	  	  
Typical Research Scale	10-50 participants, a couple of conditions	Hundreds of participants

12/50

Privacy Research in Location Sharing

- Influencing Factors
 - e.g., utility, relationship, purpose...
- Location Presentation
 - e.g., map, place names, abstractions...
- Controls and Feedback
 - e.g., black/white list, time-based rule, plausible deniability, request history...
- Location Privacy Preferences Learning
 - e.g., place naming method prediction, default policy generation...

13/50

Modeling People's Place Naming Preferences

- Objective:
 - Understand how users modulate their location information in sharing
 - E.g., sharing "at work" vs. 5000 Forbes Ave, Starbucks vs. downtown, etc.
- In study:
 - **Day Reconstruction Method**
 - Captured 26 participants' location traces for two weeks
 - Asked for their sharing preferences at each place they visited during that day

J. Lin, *et al.*, "Modeling people's place naming preferences in location sharing," In Proc. *UbiComp*, 2010.

14/50

Results

- Proposed a taxonomy of place naming methods
 - Semantic vs. Geographic, e.g. “at work” vs. “5000 Forbes ave”
 - Granularity, e.g., “Forbes & Craig”, vs. “Pittsburgh”
- Four (context-dependent) factors impact users’ choices
 - Relationship, familiarity, privacy concerns, place entropy
- Users’ preferences are predictable
 - Top category accuracy 93%
 - Granularity accuracy 89%

15/50

The Other Two Works

- K. P. Tang, J. Lin, *et al.*, "Rethinking location sharing: exploring the implications of social-driven vs. purpose-driven location sharing," In Proc. *UbiComp*, 2010.
- J. Lin, *et al.*, "A Comparative Study of Location-sharing Privacy Preferences in the U.S. and China," *PUC*, vol. under review, 2011

16/50

Lessons Learned & How They Link to the Proposed Work

- A typical Ubicomp way of collecting user preferences.
 - **Require a more scalable method.**
- Complex tradeoffs users make between utility and privacy
 - **Study similar tradeoffs in the broader context.**
- Users' privacy preferences are dynamic and complex, yet predictable
 - **Extend to a more complex problem space.**

17/50

Outline

- Background
- Related Work in Location Sharing
- **Related Work in Mobile App Privacy**
- Proposed Work: Investigating Users' Mobile App Privacy Preferences
- Schedule and Summary

18/50

Past Work in App Analysis

	Permission Analysis	Static Analysis	Dynamic Analysis
Example	[Enck, 09] [Barrera, 10] [Felt & Greenwood, 11] [Felt & Chin, 11] [Vidas, 11]	[Egele, 11] [Chin, 11] [Felt & Wang, 11] [Enck, 11] App Profiles	[Thurm, 11] : WSJ [Enck, 10]: TaintDroid [Beresford, 11] [Zhou, 11] [Hornyack, 11]
Features	Identify vulnerabilities and anomalies by analyzing the permission usage pattern	Profile apps by scanning the source codes or binary files	Capture the data flow while interacting with the apps
<p>App analysis cannot (directly) tell:</p> <ul style="list-style-type: none"> • The intention of certain behavior • How users feel about certain behavior 			

19/50

Past Work in Privacy Extensions

	Rule-Based Approach	Faking Sensitive Info
Example	[Bugiel, 11] TrustDroid [Felt & Wang, 11] Propose IPC inspection [Nauman, 10] Apex [Jeon, 12] Dr Droid & Mr. Hide	[Beresford, 11] MockDroid [Zhou, 11] TISSA [Hornyack, 11] AppFence
Features	Users can define rule based on context	Substitute fake information in the data flow
<p>?</p> <ul style="list-style-type: none"> • Can users fully understand what is necessary? • Can users correctly configure these settings? • Will these details just overwhelm the users? 		

20/50

Proposed Work

Preliminary Results

Purposes & Expectations: Understanding Users' Mental Models of Mobile App Privacy (UbiComp 2012)

Problem and Solution

■ Problem

- Existing permission screens do not help users make good trust decisions [Felt'12], [Kelley'12]
 - Few people read it → not the main task
 - People don't understand the implications
- Can we ask "others" to "digest" for users?

■ Solution

- Other human intelligence → crowdsourcing
- But what to crowdsource here?
 - Our idea: expectations and misconceptions

22/50

“Privacy as Expectations”



- Apply the idea of mental models for privacy
 - Compare what people expect an app to do vs what an app actually does
 - Emphasize the biggest gaps, the misconceptions that most people had

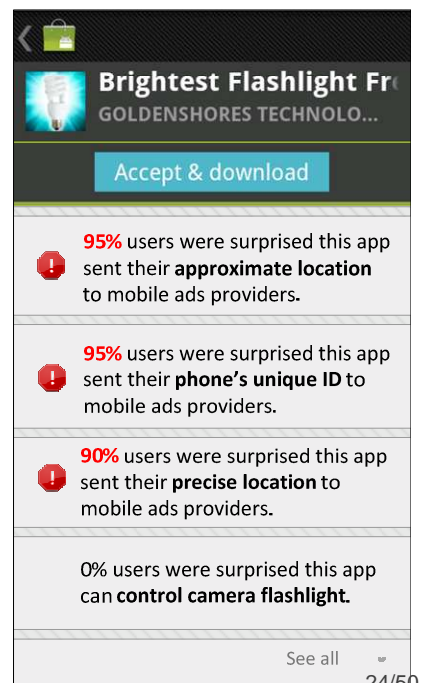

App Behavior

(What an app actually does)

User Expectation

(What people think the app does)

23/50



24/50

Our Study

- Showed crowd workers (on Mturk) screenshots and description of app (from market)
- Showed permissions one at a time
 - Only those related to privacy
- *Expectation Condition*
 - Whether they expect the app uses permission
 - Why they think the app uses permission
 - How comfortable they were with it
- *Purpose Condition*
 - We gave an explanation (based on analysis)
 - Asked how comfortable they were with it

25/50

Our Study

- Participants
 - Recruited from Mturk, US Android users only
 - Between-subjects (one condition only)
- Method
 - The top 100 popular apps in Android Market captured on Sept 12, 2011
 - Targeted types of resource
 - Location: GPS (24) and Network location (29) , Unique ID(56), Contact List (25) --- 134 app-resource pairs
 - 20 participants per pair per condition

26/50

Results of Expectation Condition


Apps Use GPS	Comfort Level (-2 – 2)	Expectation (0-100%)
Google Maps	1.52	100%
GasBuddy	1.47	95%
Weather Channel	1.45	95%
TuneIn Radio	0.60	65%
Evernote	0.15	55%
Angry Birds	-0.70	20%
Brightest Flashlight Free	-0.95	10%
Toss It	-0.95	5%

N=20 per app
GPS location data
as example

Strong correlation between expectation and perceived level of comfort, $r=0.91$

27/50

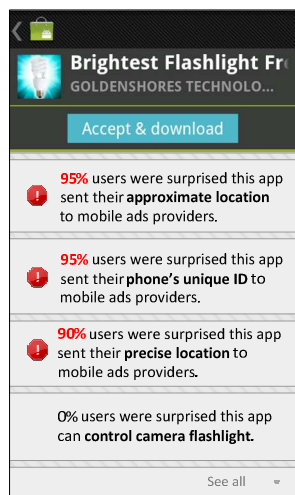
Comparing Two Conditions



- Clarifying purposes lower concerns
 - All differences statistically significant
 - Big increases for dictionary, Shazam, and others (> 1.0)

28/50

New Privacy Interface



• Key features

- Common misconceptions (expectations)
- Purpose(s) (explanations)

• Other design principles

- Simplified terms
- Only show permissions that affect privacy
- Prioritized list
- Highlight suspicious items

• Compared with existing Android permission screens

- Higher privacy awareness
- Better comprehensibility
- Required slightly less time to read

29/50

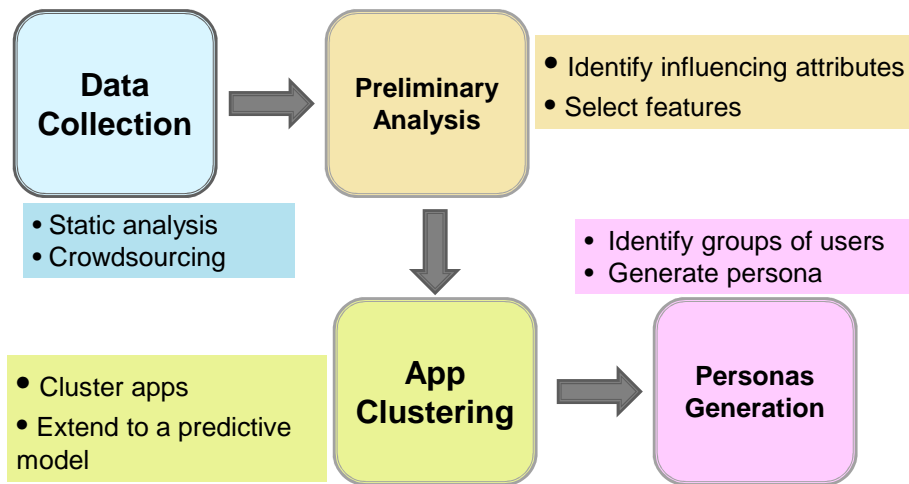
Lessons Learned...But...

- Demonstrated the feasibility of crowdsourcing
 - However, still not scalable enough...
- Identified two key factors --- expectation and purpose
 - How about other (context-free) factors?
- Proposed a preliminary design of new privacy summary interface
 - Making decision for every app, hmm...What about privacy settings?

And how? User-oriented Machine Learning

30/50

Remaining Proposed Work



31/50

Step 1: Data Collection

Objective:

Compile a dataset that include both app behaviors, and how users feel about these apps.

1. App selection -- **meta data**
2. App analysis -- **Resource usage data**
3. Crowdsourcing -- **user feedback**

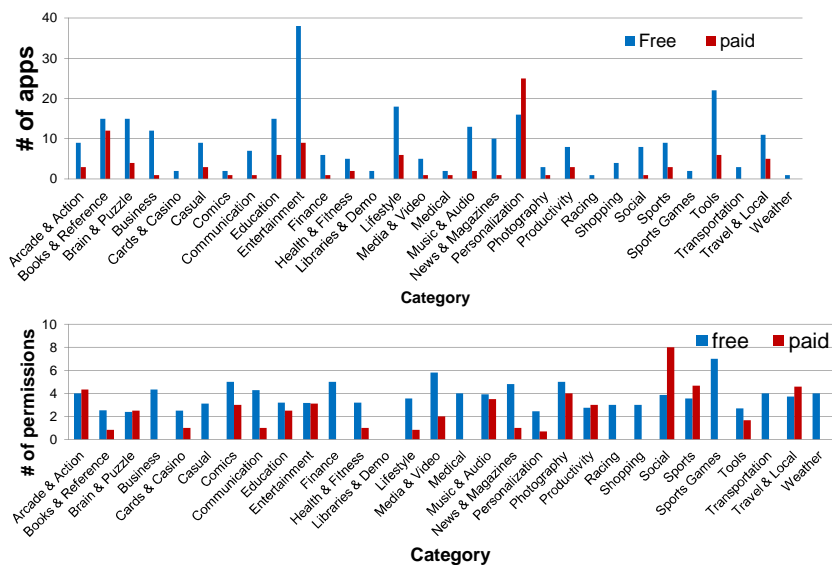
32/50

App Selection Criteria

- Representative and unbiased
 - Strategy: apps from all 30 categories, proportional to the ratio in the Market
- Feasibility of analysis
 - Strategy: filter out Non-English apps, limit the total number to ~400

33/50

Resulting App Pool



34/50

App Analysis & Crowdsourcing

- App analysis
 - Tools: apktool, app analyzer, ded
 - Resource usage data: permissions, purpose, destination URLs, 3rd party APIs, etc.
- Crowdsourcing user feedback
 - Show app screenshot, descriptions and other meta data
 - Show resource usage one at a time with purpose (10 privacy-related resources)
 - Collect participants' expectations and levels of comfort
- Collect info about participants
 - e.g., phone version, years of use, demographic information, etc.

35/50

Step 2: Preliminary Analysis

- **Objective:**
To understand how different factors impact users' decisions
- Perform regression on user feedback
 - Determine how different types of sensitive resources weigh in users' mental models
- Perform feature selection based on
 - Correlation
 - Predictability

36/50

Step 3: App Clustering

- **Objective:**
 - Identify collections of apps that elicit distinct privacy preferences
 - Extract high-level knowledge
 - E.g., Apps frequently send location info for ads concerns users,
 - E.g. , Apps communicate to only one server are more accepted
 - By-product: predictive model to estimate user acceptance

37/50

Step 3: App Clustering

- Utilize easy-to-interpret clustering algorithms
 - Choices of clustering algorithm:
 - e.g. K-means, Bottom-up
 - Other advanced algorithms (only if the aboves not working)
 - Choices of distance measure:
 - Hamming distance or weighted hamming distance
 - Euclidean distance (require proper coding of categorical data)
 - Other advance distance functions: regular simplex, symbolic covariance, etc.

38/50

Step 4: Privacy Personas Generation

- **Opportunities:**

(1) Individual differences (2) User burden

- Assuming one day

- Google redesign the privacy settings, or
- Privacy extensions developed by other parties

- What are the “right” settings?

- Effective: capture users’ need
- Usable: with low user burden

- **Objective:** Can we identify a set of **GOOD** default settings

- → **Privacy Personas**



39/50

Step 4: Privacy Personas Generation

- **Opportunities:**

(1) Individual differences (2) User burden

Privacy Persona is


A set of privacy policy shared by a group of users, e.g.

“I am willing to disclose my location for the functionality purpose.

I am NOT willing to disclose my location for advertising or my call log for any purpose.”

Objective: Can we identify a set of **GOOD** default settings

- → **Privacy Personas**



40/50

Step 4: Privacy Personas Generation

- How to generate default privacy personas?
 1. Learn the privacy policy for each participant -> (preference vectors)
 2. Identify groups of users share similar preferences (clustering)
 3. Extract average policy of each group (cluster center)
- Evaluate generated personas
 - Repeat crowdsourcing for new apps and new participants
 - Lab studies (if time allows)

41/50

Summary

- Complement existing mobile privacy research with in-depth user research
- A new way of looking at privacy, i.e. "Privacy as Expectation"
- A valuable dataset containing both app behavioral attributes and users privacy ratings
- Clusters of mobile apps that elicit distinct privacy concerns
- A set of privacy personas that can simplify privacy settings

42/50

How This Work Benefit Diff Parties

- Provide **mobile app markets** models to evaluate apps from privacy perspectives
- Inform the design usable and efficient privacy interfaces and settings of existing **mobile OS**
- Help **developers** to understand the privacy implications of their apps

43/50

Schedule



Aim at finishing by Aug 2013

Publication opportunities:

MobiSys'13, submission deadline Dec 2012

UbiComp'13, submission deadline March/ April 2013

44/50

Acknowledgement

- This work is supported by **CyLab** at Carnegie Mellon under grants DAAD19-02-1-0389 and W911NF-09-1-0273 from the **Army Research Office** and by **Google**. Support was also provided by the **National Science Foundation** under Grants CNS-1012763 and CNS-0905562.

45/50

Selected References -1

- App Analysis
 - S. Thurm and Y. I. Kane, "Your Apps are Watching You," *WSJ*, 2011.
 - E. Chin, *et al.*, "Analyzing inter-application communication in Android," In Proc. *MobiSys*, 2011
 - M. Egele, *et al.*, "PiOS: Detecting Privacy Leaks in iOS Applications," In Proc. *NDSS*, 2011.
 - W. Enck, *et al.*, "TaintDroid: An Information-Flow Tracking System for Realtime Privacy Monitoring on Smartphones," In Proc. *OSDI* 2010.
 - Y. Zhou, *et al.*, "Taming Information-Stealing Smartphone Applications (on Android)," In Proc. *TRUST*, 2011
- Security Extensions
 - A. Beresford, *et al.*, "MockDroid: trading privacy for application functionality on smartphones," In Proc. *HotMobile*, 2011.
 - P. Hornyack, *et al.*, "These aren't the droids you're looking for: retrofitting android to protect data from imperious applications," In Proc. *CCS*, 2011.
 - J. Jeon, *et al.*, "Dr. Android and Mr. Hide: Fine-grained security policies on unmodified Android," *Working paper*, available at <http://www.cs.ucla.edu/~jeff/docs/drandroid.pdf>, 2012.
- Android Permission Studies
 - A. P. Felt, *et al.*, "Android Permissions: User Attention, Comprehension, and Behavior," UCB/EECS-2012-26, University of California, Berkeley, 2012
 - P. G. Kelley, *et al.*, "A Conundrum of permissions: Installing Applications on an Android Smartphone," In Proc. *USEC*, 2012.
 - A. P. Felt, *et al.*, "I've Got 99 Problems, But Vibration Ain't One: A Survey of Smartphone Users' Concerns," In Proc. *SOUPS*, 2012.

46/50

Selected References -2

■ Location Sharing

- B. Brown, *et al.*, "Locating Family Values: A Field Trial of the Whereabouts Clock," In Proc. *UbiComp*, 2007.
- G. Iachello, *et al.*, "Control, Deception, and Communication: Evaluating the Deployment of a Location-Enhanced," ed: unknown, 2008.
- M. Benisch, *et al.*, "Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs," *Personal and Ubiquitous Computing*, 2010.
- S. Consolvo, *et al.*, "Location disclosure to social relations: why, when, & what people want to share," In Proc. *CHI*, 2005.
- J. Cranshaw, *et al.*, "User-Controllable Learning of Location Privacy Policies with Gaussian Mixture Models," In Proc. *AAAI*, 2011.
- J. Y. Tsai, *et al.*, "Location-Sharing Technologies: Privacy Risks and Controls," In Proc. *TPRC*, 2009.

47/50

Q & A

■ Jiali Lin
jialiul@cs.cmu.edu
School of Computer Science
Carnegie Mellon University


48/50

Backup Slides

Mobile Apps Provide Attractive Services— shorten early slides



You were observed at this location
from 15:35 Aug 12 (Wed) to 16:24 Aug 12 (Wed)

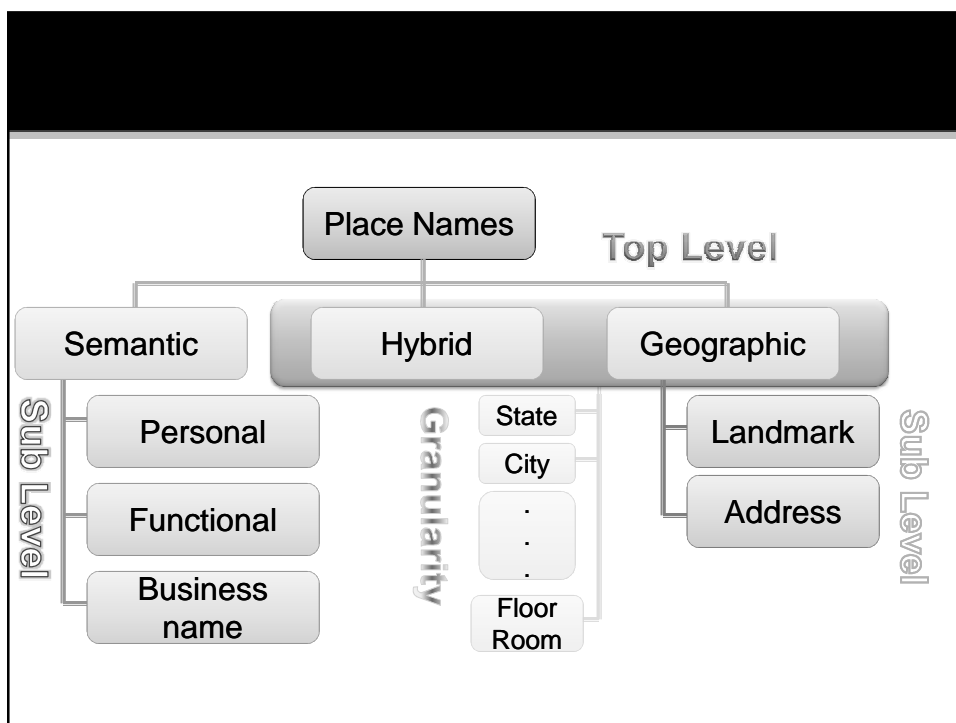


they visited.


- Questions were asked for four social groups
 - family member,
 - close friend,
 - acquaintance,
 - stranger.

Imagine that **Mary** (your family member) wanted to know where you were at the given time period.

1. How comfortable would you be to let her know where you were at this time?
(1: not comfortable at all, 7: extremely comfortable)
2. How familiar is **Mary** with this location?
(1: not familiar at all, 7: extremely familiar)
3. What terms or phrases (place name) would you use to refer to this location if you want to tell her where you were?
(e.g.: office, 1211 Hamburger Hall, Starbucks, grocery store, etc.)



App Name: Toss it



Toss a ball of crumpled paper through a hoop. MILLIONS of Android games on the market -- FREE!
 - Simple yet challenging game
 - Challenge your friends to toss
 - Toss that paper through a hoop
 And if you like Toss It, check out Toe LIVE! - aiMinesweeper (Four)

1. Have you used this app?
☐ Yes ☐ No

2. What category do you use it for?
 (required)
☐ Game ☐ Other

The Expectation Condition
 Please provide any comments of this app you may have below.

3. Suppose you have installed Toss it on your Android device, would you expect it to access your **precise location**? (required)
☐ Yes ☐ No

Toss it does access users' **precise location information**.

4. Could you think of any reason(s) why this app would need to access this information? (required)
 = precise location is necessary for this app to serve its purpose

Based on our analysis, Toss it accesses user's **precise location information** for **targeted advertising**.

3. Suppose you have installed Toss it on your Android device, do you feel comfortable letting it access your **precise location**? (required)
☐ Very comfortable
☐ Somewhat comfortable
☐ Somewhat uncomfortable
☐ Very uncomfortable

☐ Somewhat comfortable
☐ Somewhat uncomfortable
☐ Very uncomfortable

The Most Unexpected Resource Usages

Resource	App name	% Expected	Avg Comfort
Network Location	Brightest Flashlight	5%	-1.25
	Toss It	10%	-1.15
	Angry Birds	10%	-0.43
	Air Control Lite	20%	-0.55
	Horoscope	20%	-1.05
GPS Location	Brightest Flashlight	10%	-0.95
	Toss It	5%	-0.95
	Shazam	20%	-0.05
Device ID	Brightest Flashlight	5%	-1.35
	TalkingTom Free	10%	-0.78
	Mouse Trap	15%	-0.85
	Dictionary	15%	-0.69
	Tiny Flashlight	20%	-0.80
	Ant Smasher	20%	-1.13
	FxCamera	20%	-0.73
Contact List	Horoscope	20%	-1.03
	Backgrounds HD	10%	-1.35
	Wallpapers	20%	-0.70
	Pandora	20%	-0.70
	GO Launcher EX	20%	-0.75

The comfort rating was ranging from -2.0 (very uncomfortable) to +2.0 (very comfortable).

- strong correlation observed ($r=0.91$) between people's expectation and their subjective feelings
- Perceived necessity guide users to make trust decisions or prompted them to take different actions.
- W27 "Why does a flashlight need to know my location? I love this app, but now I know it access my location, I may delete it." (Brightest Flashlight)
- W56 "I do not feel that games should ever need access to your location. I will never download this game." (Toss it)

The Quality of Crowdsourced Data

- To prevent gaming of our study
 - Crowd workers' lifetime approval rate >75%
 - Limit to Android users
 - Quality control question
- Similar results comparing to lab study
 - Mean Square Errors are negligible

MSE	Network Loc	GPS loc	Contact List	Unique ID
expectation [0,1]	0.0354	0.0303	0.0353	0.0363
comfort level [-2,+2]	0.7081	0.8136	0.6749	0.3067

Lay Users Have a Hard Time Identifying the Purposes of

Resource Type	Resource used for [1] Major functionality [2] Tagging or sharing [3] Advertising or market analysis	% of accurate guess	% of no idea
Contact List (25)	[1]-----20	56%	8%
	[2]-----2	28%	35%
	[1]+[2]-----2	19%	16%
	[1]+[2]+[3]-----1	27%	14%
GPS Location (24)	[1]-----14	74%	11%
	[2]-----4	80%	10%
	[3]-----2	35%	55%
	[1]+[3]-----3	15%	27%
Network Location (29)	[2]+[3]-----1	15%	40%
	[1]-----15	77%	8%
	[2]-----2	55%	10%
	[3]-----7	29%	63%
Device ID (56)	[1]+[3]-----3	15%	22%
	[2]+[3]-----2	13%	25%
	[1]-----14	51%	29%
	[3]-----30	22%	58%
	[1]+[3]-----12	7%	55%

- TaintDroid was used to analyze the ground truth.
- We manually categorized apps into 3 categories:
 - For major functionality
 - for sharing and tagging (or supporting other minor functions)
 - for target advertising or market analysis
- Accuracy never exceeded 80% even for purely functionality purposes.
- Very low accuracy when sensitive resources used for multiple purposes

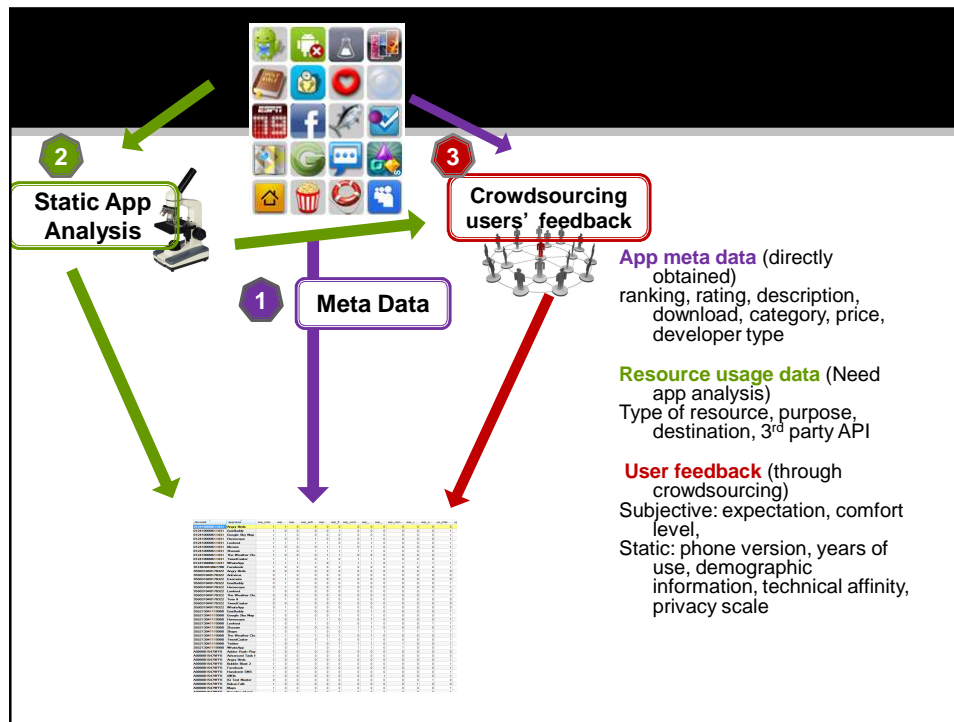
Clarifying the Purpose May Ease Worries

Resource Type	comfort rating w/ purpose	comfort rating w/o purpose	df	T	p
Device ID	0.47(0.30)	-0.10(0.41)	55	7.42	0.0001
Contact List	0.66(0.22)	0.16(0.54)	24	4.47	0.0002
Network Location	0.90(0.53)	0.65(0.55)	28	3.14	0.004
GPS Location	0.72(0.62)	0.35(0.73)	23	3.60	0.001

The comfort rating was ranging from -2.0 (very uncomfortable to +2.0 (very comfortable).

Category	% in market	Paid ratio	#paid	#free	Category	% in market	Paid ratio	#paid	#free
Android Application					Sports	3.2%	24%	3	9
Books & Reference	7.4%	47%	13	15	Tools	7.4%	23%	6	22
Business	3.5%	9%	1	12	Transportation	1.0%	17%	0	3
Comics	1.0%	43%	1	2	Travel & Local	4.3%	33%	5	11
Communication	2.2%	19%	1	7	Weather	0.4%	17%	0	1
Education	5.6%	30%	6	15	Libraries & Demo	0.6%	10%	0	2
Entertainment	12.0%	20%	9	38	Total applications	86.8%		89	235
Finance	1.9%	15%	1	6	Android Games				
Health & Fitness	2.2%	32%	2	5	Arcade & Action	3.3%	25%	3	9
Lifestyle	6.3%	27%	6	18	Brain & Puzzle	5.0%	21%	4	15
Media & Video	1.8%	21%	1	5	Cards & Casino	0.9%	27%	0	2
Medical	1.0%	40%	1	2	Casual	3.2%	24%	3	9
Music & Audio	3.9%	13%	2	13	Sports Games	0.7%	26%	0	2
News & Magazines	2.9%	9%	1	10	Racing	0.4%	15%	0	1
Personalization	10.6%	60%	25	16	Total games	13.4%		10	38
Photography	1.3%	25%	1	3					
Productivity	2.9%	27%	3	8					
Shopping	1.2%	9%	0	4					
Social	2.3%	11%	1	8					

* <http://www.appbrain.com/stats/android-market-app-categories>



Step 4: Privacy Personas Generation

- 1: Learn the privacy policy for each user
 - Rearrange per-app preferences into a rule-based policy
 - Encode policy into a vector, each entry represents preference of an app cluster
- 2: Perform clustering on these vectors
 - Group users who have similar policy vectors
 - E.g., K-means with Hamming distance
- 3: Learn a default persona for each cluster
 - Find the center of each cluster by averaging the policy vectors within each cluster