USABLE PRIVACY POLICY PROJECT

Scaling Privacy in the Age of the Internet of Things: Could AI Hold the Solution?

Norman Sadeh

Carnegie Mellon University

Carnegie Mellon University

CLIP Center on Law and Information Policy AT FORDHAM LAW SCHOOL

CIS The Center for Internet and Society

Outline

- Privacy in the age of IoT
- Notice and Choice is Broken
- The Usable Privacy Policy Project
- Can We Crowdsource Policy Annotations?
- How Can Machine Learning and Natural Language Processing Help?
- Nudging Users to Review their Privacy Settings
- Next Steps

The Internet of Things

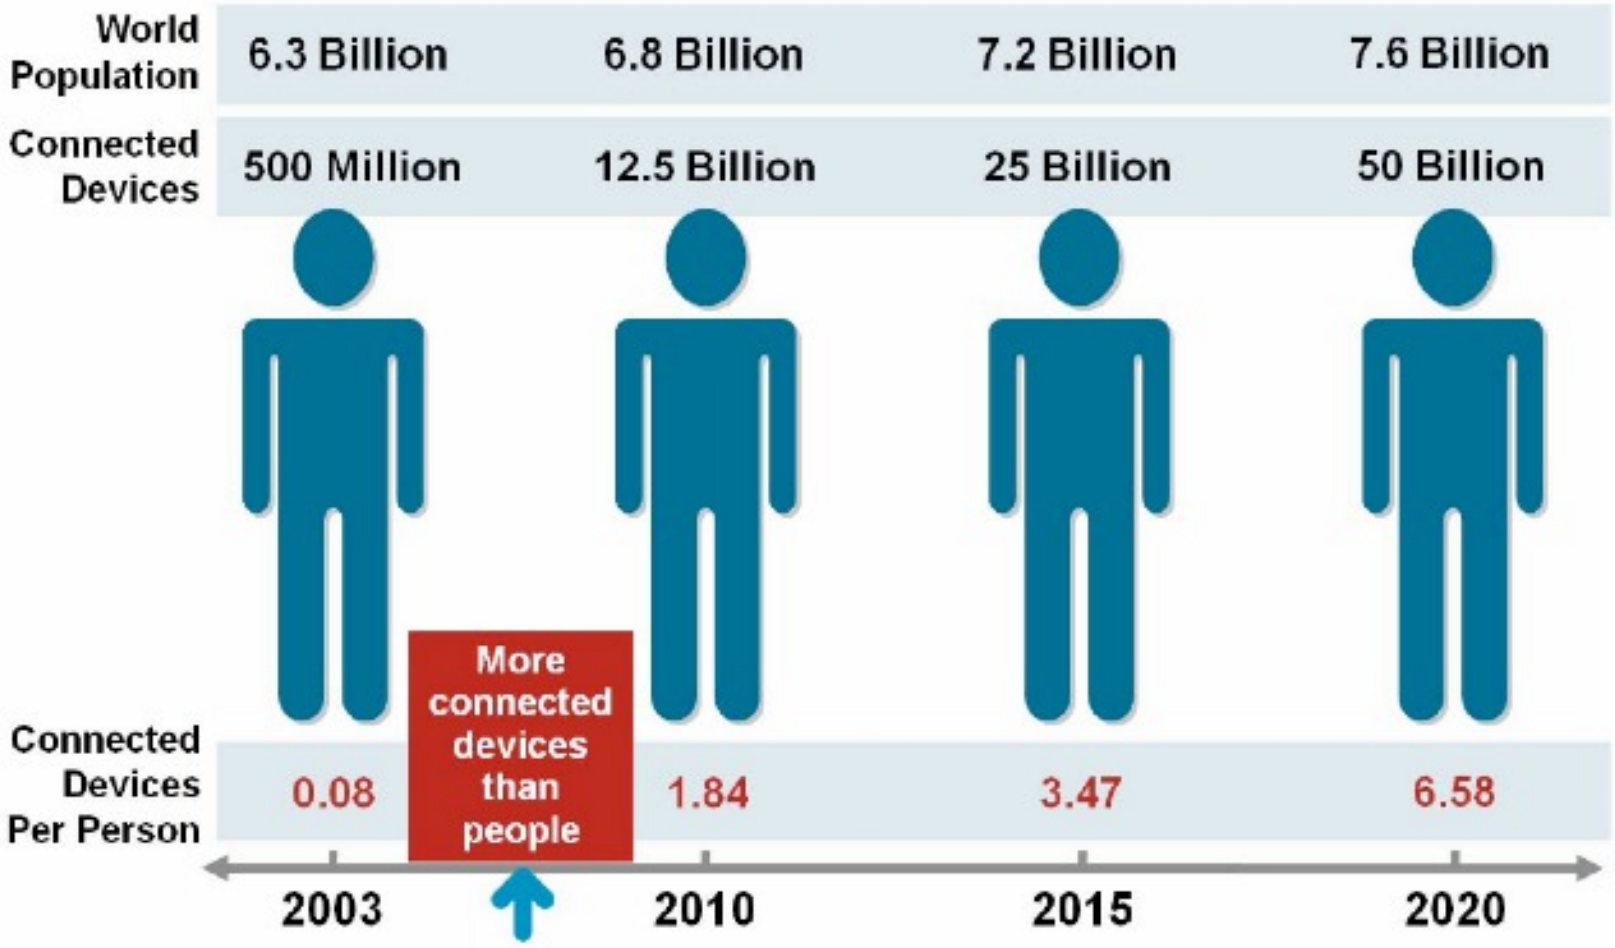

Source: Cisco IBSG, April 2011

Privacy Implications

- As we go about our daily lives, we interact with a number of devices, applications and services
- **Many of these devices applications and services may collect, share and mine data about us**
  - Many potential benefits
  - …but also many potential risks

Benefits …and Risks

- Your home thermostat accesses your calendar to start the AC or furnace in time for when you return from work
- Your smartwatch shares your heart rate with your doctor
- …**But would you want…**
  - your phone to also report your driving habits to your car insurance provider?
  - your blood pressure to be sent to your health insurance provider?

“Notice and Choice”

- **Not everyone feels the same way** about these potential privacy risks
- “Notice and Choice” is intended to support **informed consent**
  - Enshrined in many legal documents
    - Including Hong Kong’s Personal Data (Privacy) Ordinance

Notice and Choice in Practice

- Privacy policies
  - Websites, mobile apps,
- Privacy settings
  - Smartphones, browsers, facebook

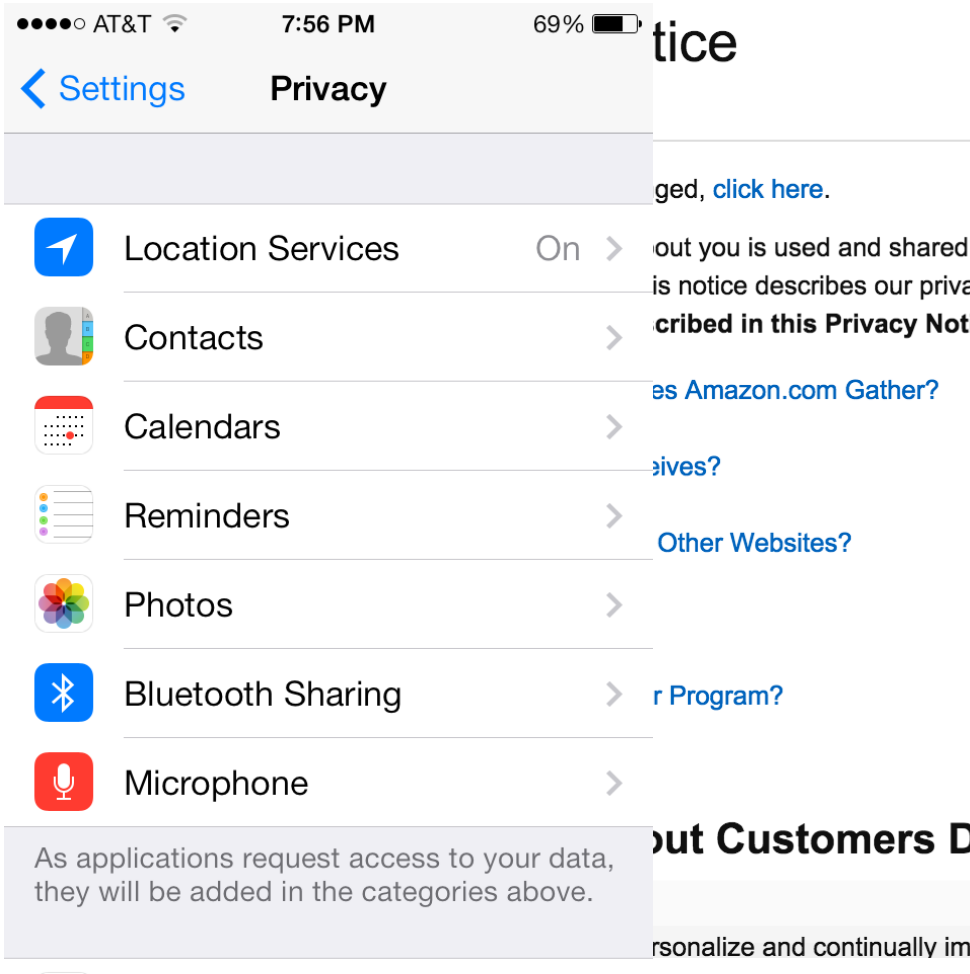

A First Quick Question

- How many of you have read a privacy policy over the past month?

People Are Feeling Helpless

- Reading a privacy policy takes about 10 minutes…or about 200 hours/year for an average Internet user… (McDonald & Cranor 2009)
- A typical mobile phone user with 40 mobile apps each requesting 3 permissions would have to configure 120 settings (Lin,Liu, Sadeh, Hong 2014)

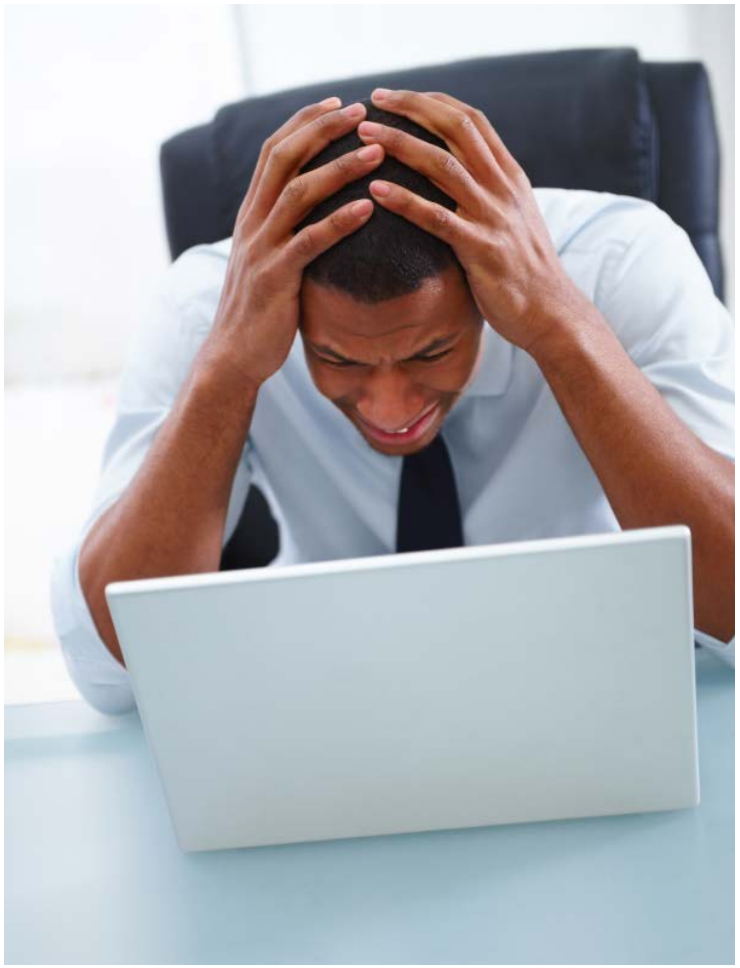

The Usable Privacy Policy Project

- Annotate natural language website privacy policies to **capture key policy considerations – those that matter to users**
- Develop **concise, intuitive and effective UIs to convey key information** to users
  - Ultimately in a **personalized** fashion

Approach

- Combining **machine learning**, **natural language processing** and **crowdsourcing** to scale policy annotation
- **Modeling people's privacy preferences** to focus on those questions users care about
- **Effective UIs** – browser plug-in(s)
- **Analysis** of website privacy policies
  - Ambiguity, compliance, stated practices
  - Across sectors, within sectors
  - Language addressing specific issues

Tightly Interconnected Threads

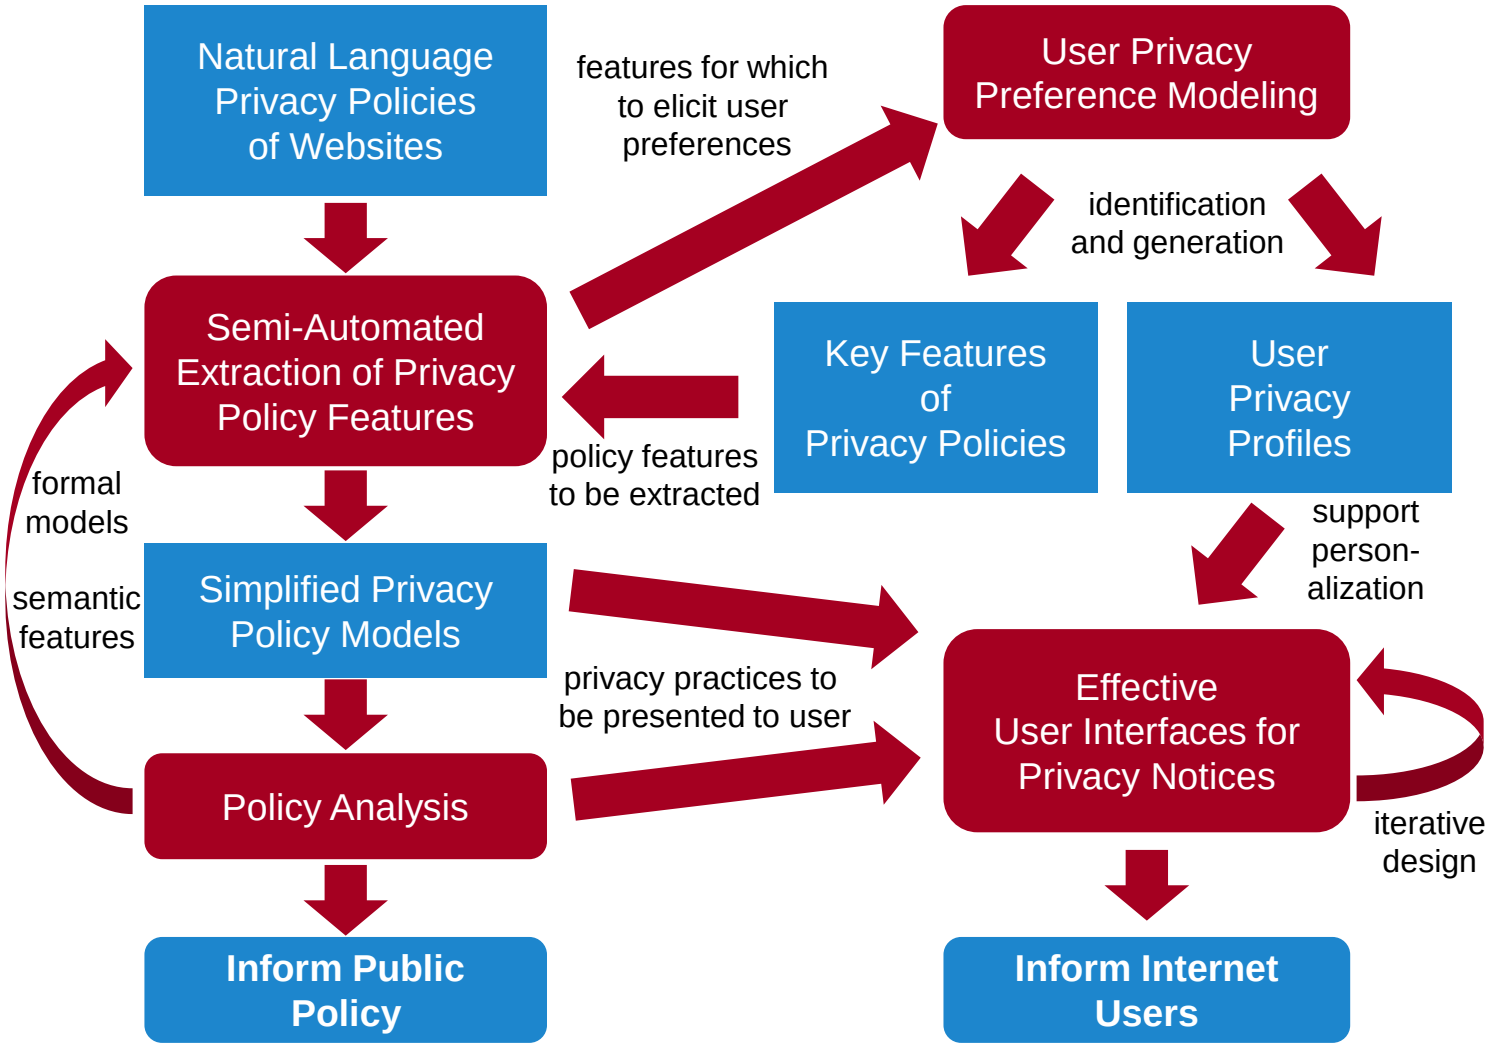

Initial Focus on 9 Questions

- **Collection** of **contact, location, health, financial** data
- **Sharing** of **contact, location, health, financial** data
  - Distinguishing between different options (e.g. sharing to support core service vs. sharing for a secondary purpose)
- **Deletion** of personal data
- Including "***Not clear***" and "***not addressed***" options

Crowdsourcing Answers to 9 Questions

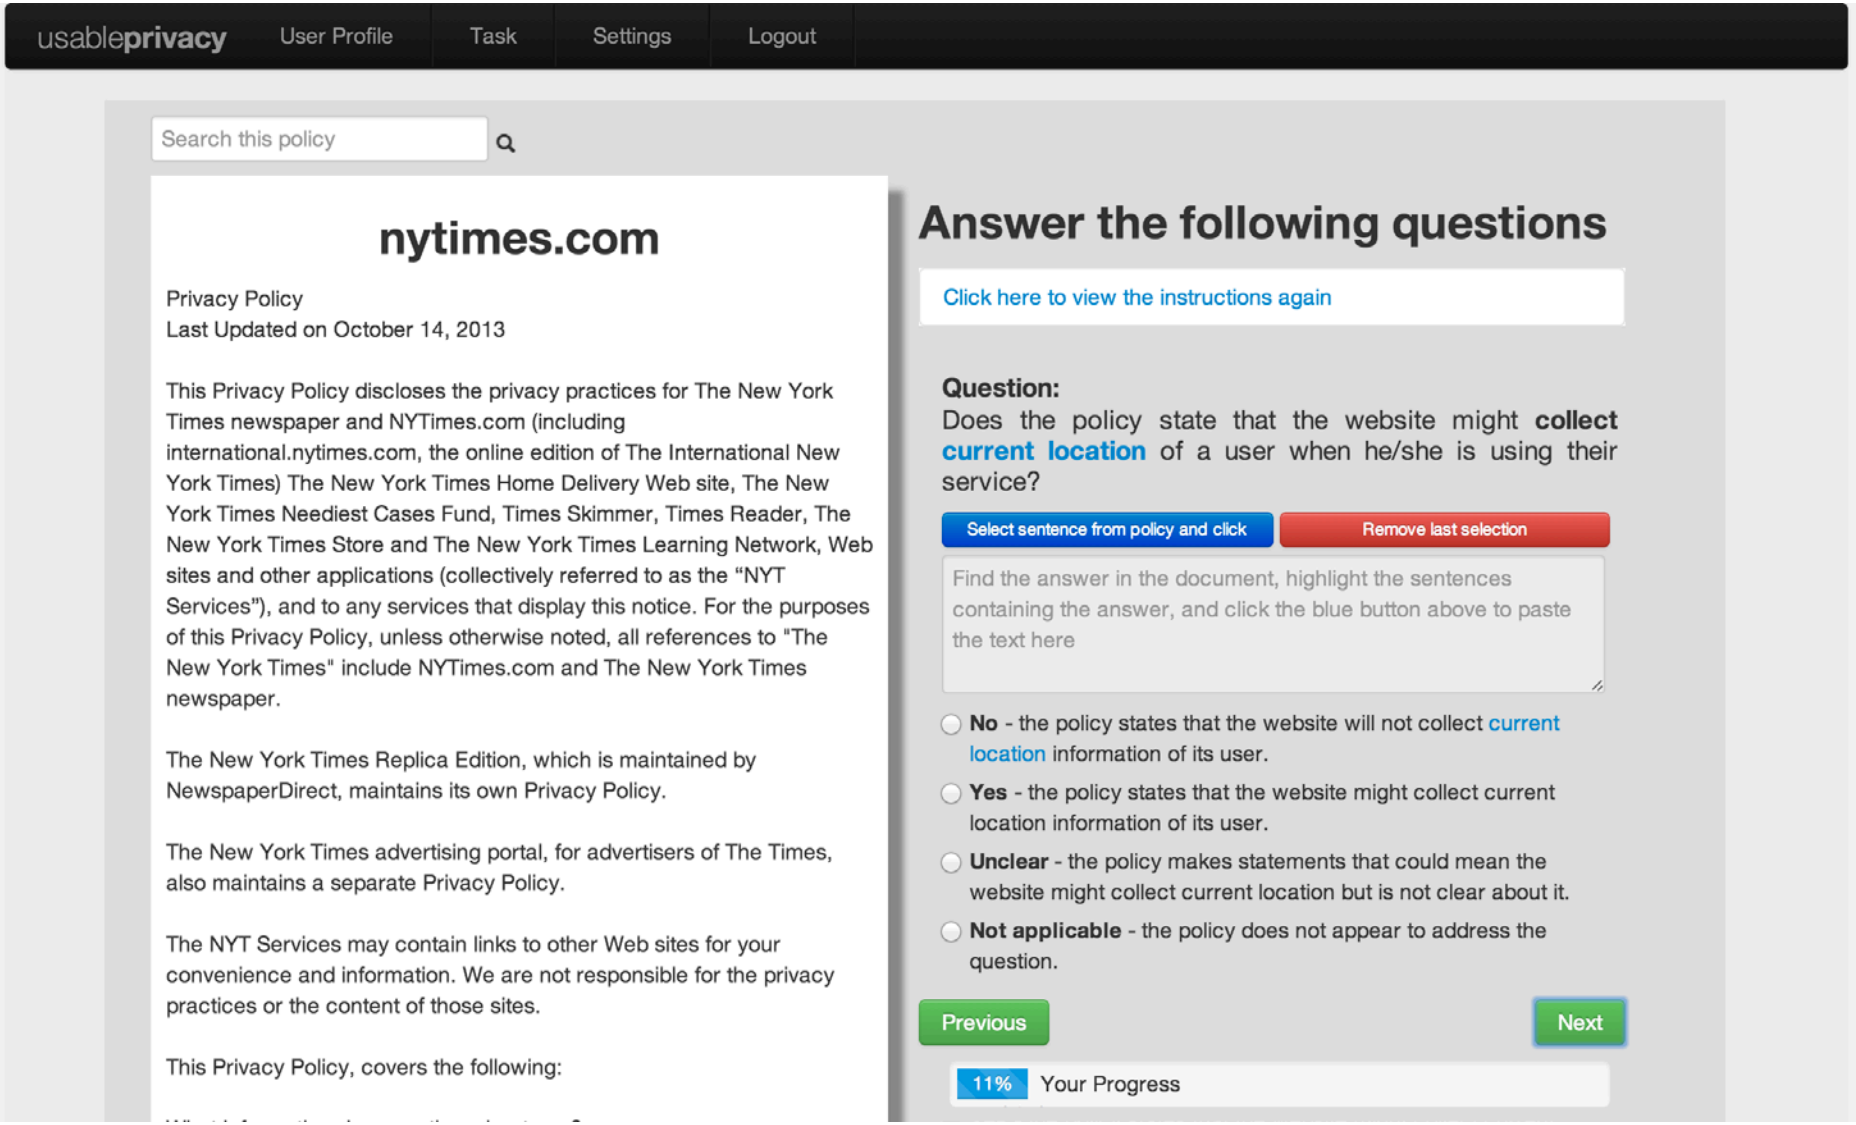

Initial Work – Amazon Mechanical Turk

amazon mechanical turk beta
Artificial Artificial Intelligence

Your Account | HITs | Qualifications

Introduction | **Dashboard** | **Status** | **Account Settings**

Mechanical Turk is a marketplace for work.

We give businesses and developers access to an on-demand, scalable workforce.
Workers select from thousands of tasks and work whenever it's convenient.

301,177 HITs available. View them now.

Make Money
by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. Find HITs now.

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work

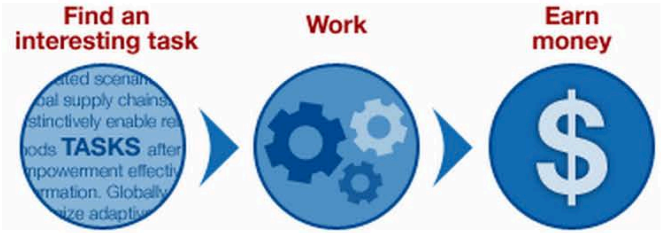

Find HITs Now

Get Results
from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. Get Started.

As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results

Get Started

Does Crowdsourcing Even Work?

Questions:

- Do Turkers converge (80%) on the correct answer?
- Gold standard obtained from skilled annotators with experience reading privacy policies

- 26 policies, 9 questions, **10 turkers per policy-question pair**

# policy-question pairs with gold standard	80% turkers agree (same answer)	80% turkers agree (different answer)	Turker agreement below 80%
198	36 (75%)	1 (2%)	11 (23%)

Crowdsourcing Does Work…But…

- A single user will often make mistakes…but collectively crowdworkers are able to accurately annotate **75% of policy-question pairs**
  - And they **rarely seem to converge on erroneous annotations**
- …But the process is **time consuming**: around 20 minutes per policy…

Alternatively Could we Fully Automate?

Question Category	Avg. Accuracy
Information collected or inferred	73%
Entities with which info may be shared	**63%**
Retention and Access	**64%**
Purposes	**59%**
Consent Model (Can users limit?)	**68%**
Choice method	74%
Security and other practices	75%

Our Approach

- **Combine crowdsourcing with machine learning and natural language processing**
- Find ways of decomposing the problem
  - Reduce the number and complexity of the tasks assigned to crowdworkers

Paragraph Sequencing

- Could we automatically organize paragraphs based on the privacy issues they discuss?
- And use this as a basis for simplifying the work of crowdworkers

Paragraph Sequencing – Problem Overview (I)

amazon.com

What About Cookies?
Cookies are unique identifiers that we transfer to your device to enable our systems to recognize your device and to provide features such as 1-Click purchasing, Recommended for You, personalized advertisements on other Web sites...

...Because cookies allow you to take advantage of some of Amazon.com's essential features, we recommend that you leave them turned on. For instance, if you block or otherwise reject our cookies, you will not be able to add items to your Shopping Cart, proceed to Checkout, or use any Amazon.com products and services that require you to Sign in...

Example of 2 paragraphs discussing cookies

Paragraph Sequencing – Problem Overview (II)

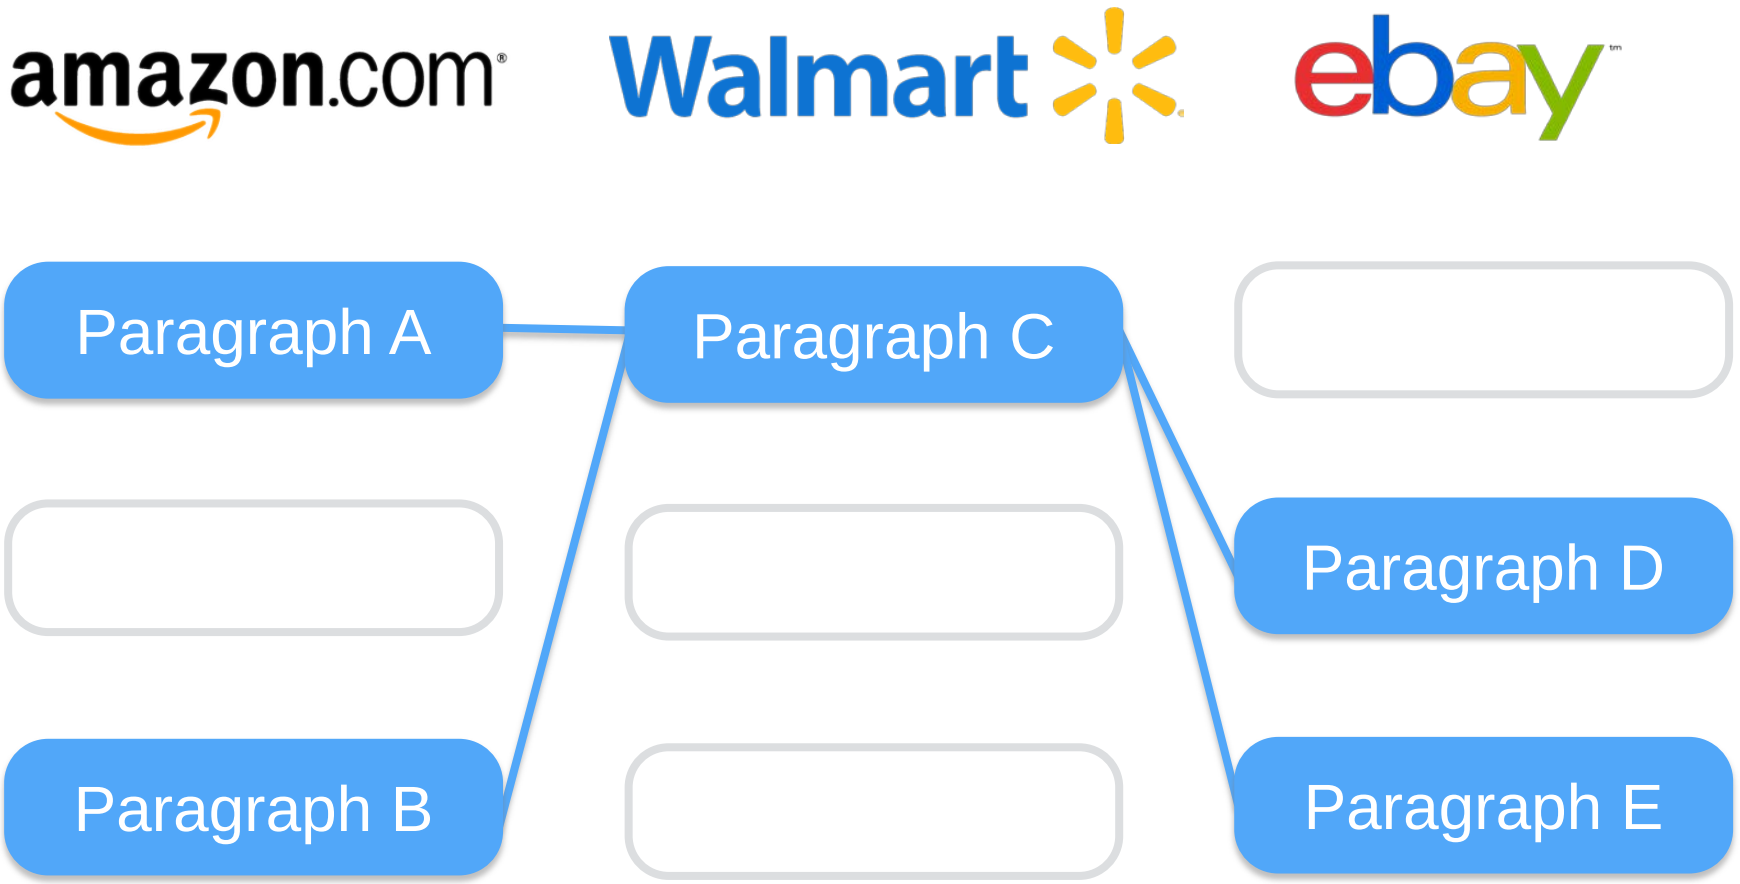

Paragraph Sequencing – Gold-standard

Instructions:

- Your task is to read `two text descriptions` extracted from `website privacy policies`. Example privacy policies include: Amazon Privacy Notice, Ebay Privacy Policy.
- You will `select 1~3 keywords` for each description, and `check "yes" or "no"` to decide if the two descriptions are discussing the same `privacy issue` or not.
- `Example privacy issues` may include: collection of personal information, sharing information with third parties, cookies and other tracking techniques, data security, policies for children, contact of the merchant, etc.

Example:

Description 1:

```
Choice/Opt-out. If we ever send you information by email concerning new products, services or information that y
ou did not expressly request, we will provide you with an email address by which you may request no further noti
ces.
```

`Keywords:` `choice, opt-out`

Description 2:

```
You may also opt-out of your participation in most of the ABC's digital services. Information about how to opt-o
ut will be provided in the particular service. However, you should be aware that the ABC may continue to store p
ersonal information provided by you prior to you opting-out.
```

`Keywords:` `opt-out`

`Are they discussing the same privacy issue?` `Yes.` Both of the descriptions are discussing the "opt-out" options provided to the website users, although the detailed terms may be different.

Privacy Policy Dataset

1010 privacy policies collected during December 2013 to January 2014, ranging over 15 website categories

Arts	Business	Computers	Games	Health
Home	Kids and Teens	News	Recreation	Reference
Regional	Science	Shopping	Society	Sports

Fei Liu, Rohan Ramanath, Norman Sadeh, Noah A. Smith. A Step Towards Usable Privacy Policy: Automatic Alignment of Privacy Statements. In Proceedings of the 25th International Conference on Computational Linguistics (COLING 2014), Dublin, Ireland, August 2014.

Paragraph Sequencing – Approach

- Hidden Markov model representation
  - Each hidden state represents a privacy issue
  - Each observation represents a text segment

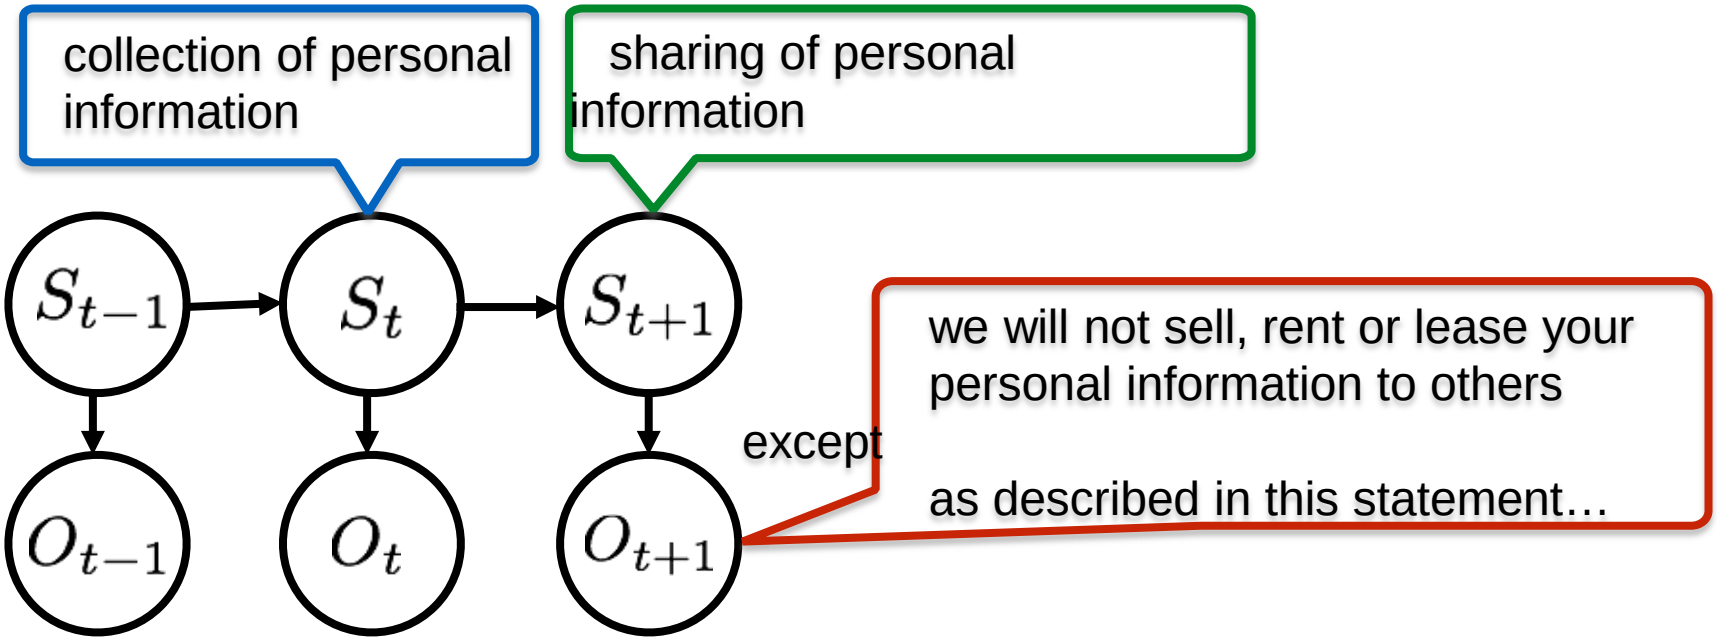

Paragraph Sequencing – Algorithms

- Hidden Markov model representation
  - Each hidden state represents a privacy issue
  - Each observation represents a text segment

- System comparison
  - CLUTO: a greedy divising clustering algorithm
  - EM-HMM: expectation maximization
  - VB-HMM: variational Bayesian inference

Paragraph Sequencing – Evaluation

- Gold-standard: human-labeled paragraph pairs

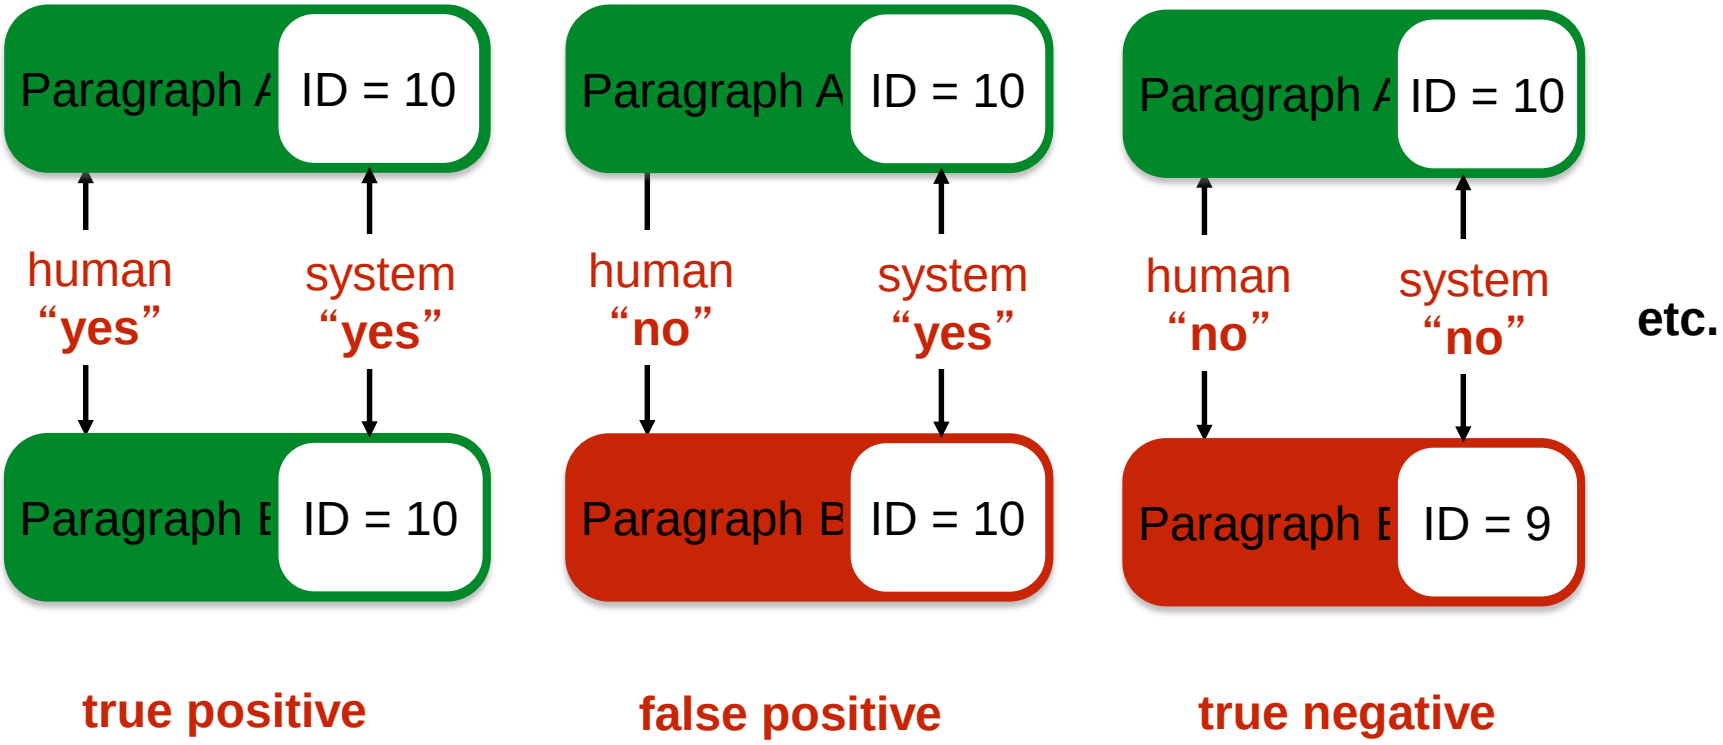

Paragraph Sequencing – Results

- Gold-standard: human-labeled paragraph pairs
- System results
  - HMMs outperform clustering when K is in the range of [5, 15)**, best performance achieved by HMMs at 87% f-score**
  - Using two levels of text granularities (paragraphs and sections), systems achieve similar results on f-scores

Highlighting Technique

- For each of the 9 questions:
  - Look for the presence of key combinations of terms in text highlighted by skilled annotators in answering these questions
  - Learn models that can be used to highlight relevant paragraphs in crowdworker UI

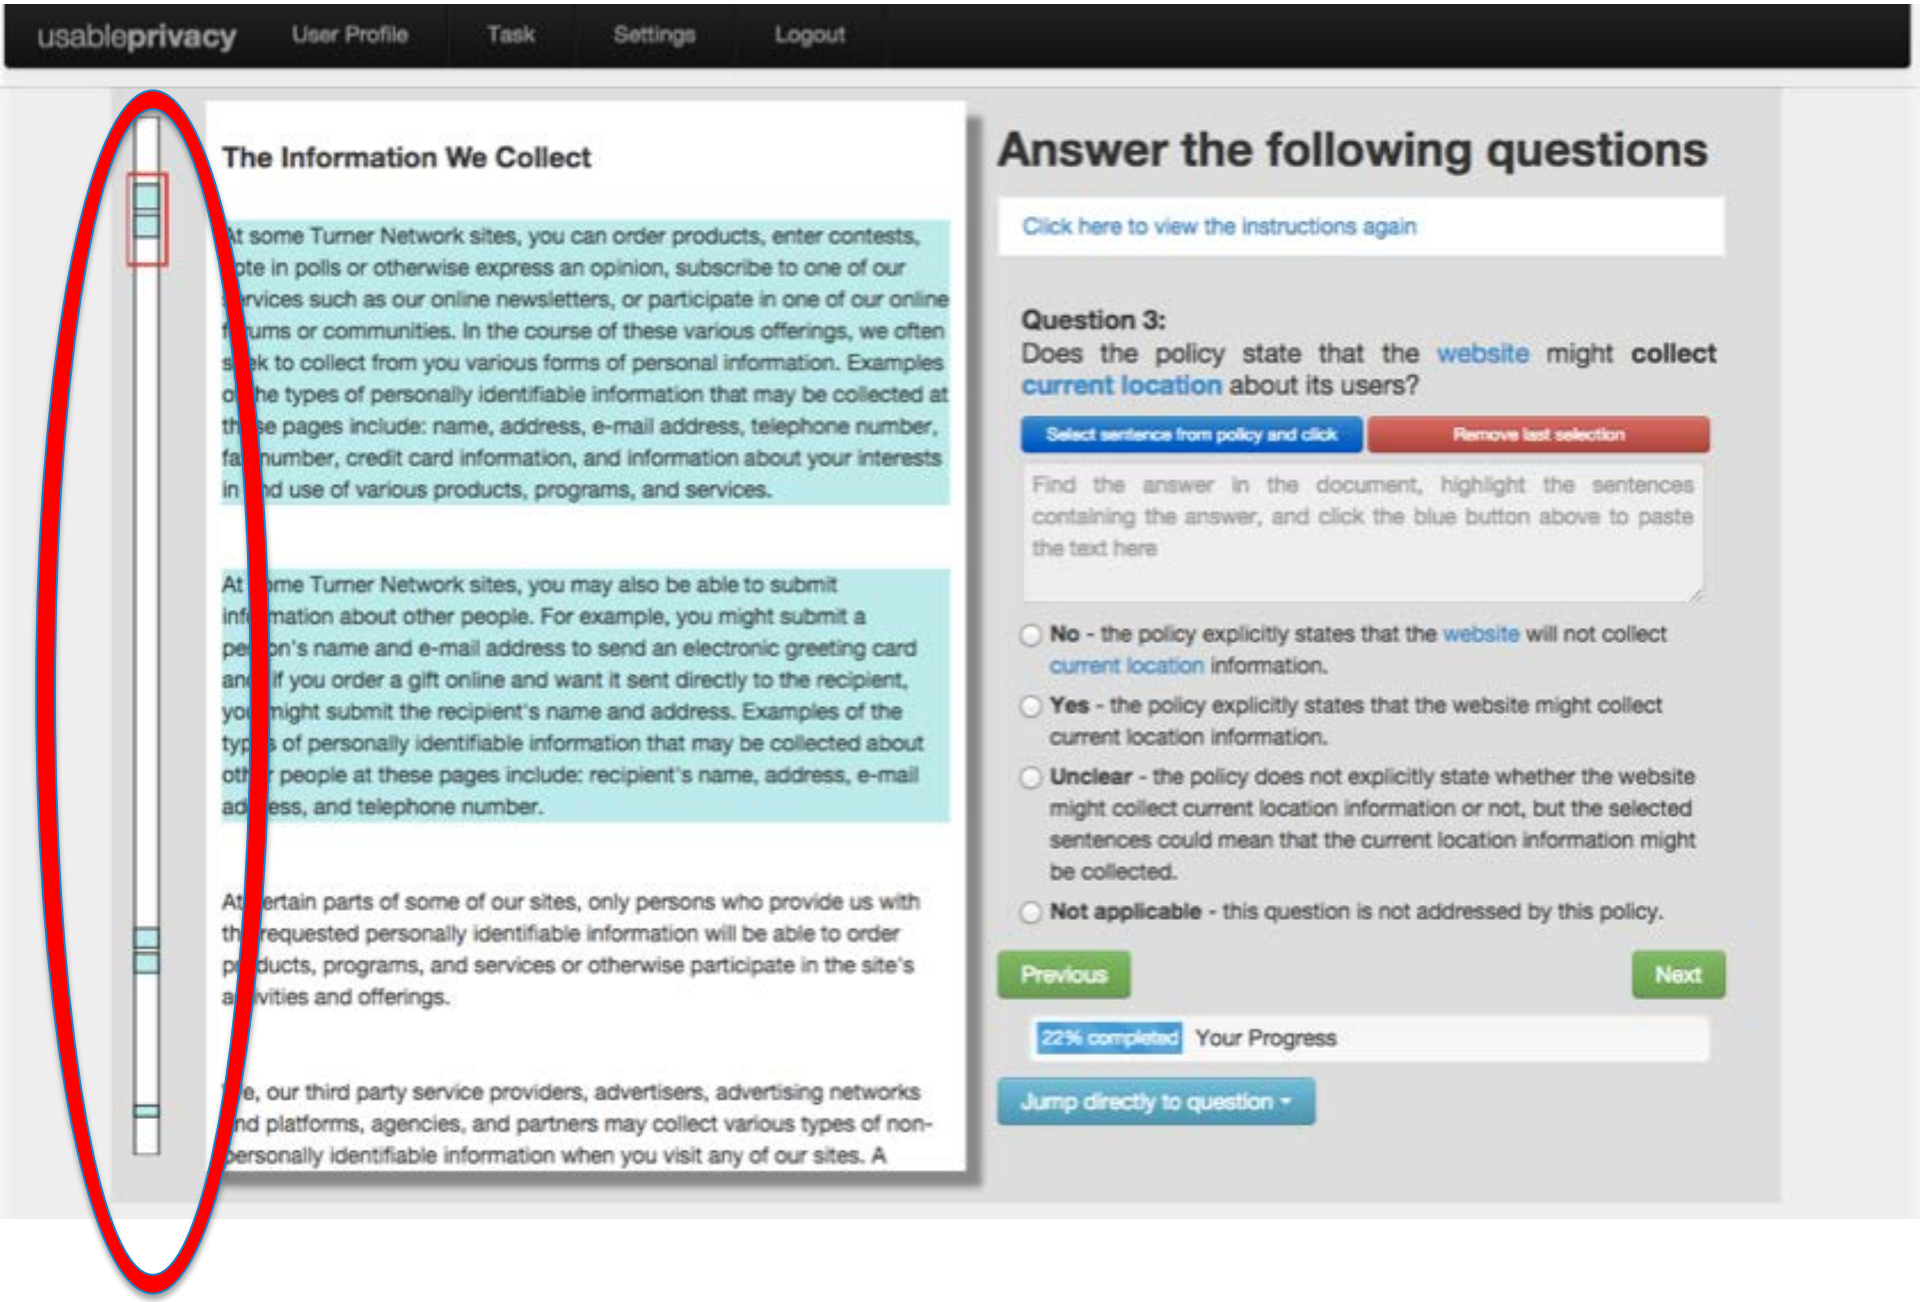
usableprivacy
User Profile
Task
Settings
Logout
The Information We Collect
At some Turner Network sites, you can order products, enter contests, vote in polls or otherwise express an opinion, subscribe to one of our services such as our online newsletters, or participate in one of our online forums or communities. In the course of these various offerings, we often seek to collect from you various forms of personal information. Examples of the types of personally identifiable information that may be collected at these pages include: name, address, e-mail address, telephone number, fax number, credit card information, and information about your interests in and use of various products, programs, and services.
At some Turner Network sites, you may also be able to submit information about other people. For example, you might submit a person's name and e-mail address to send an electronic greeting card and if you order a gift online and want it sent directly to the recipient, you might submit the recipient's name and address. Examples of the types of personally identifiable information that may be collected about other people at these pages include: recipient's name, address, e-mail address, and telephone number.
At certain parts of some of our sites, only persons who provide us with the requested personally identifiable information will be able to order products, programs, and services or otherwise participate in the site's activities and offerings.
We, our third party service providers, advertisers, advertising networks and platforms, agencies, and partners may collect various types of non-personally identifiable information when you visit any of our sites. A
Answer the following questions
Click here to view the instructions again
Question 3:
Does the policy state that the website might collect current location about its users?
Select sentence from policy and click
Remove last selection
Find the answer in the document, highlight the sentences containing the answer, and click the blue button above to paste the text here
No - the policy explicitly states that the website will not collect current location information.
Yes - the policy explicitly states that the website might collect current location information.
Unclear - the policy does not explicitly state whether the website might collect current location information or not, but the selected sentences could mean that the current location information might be collected.
Not applicable - this question is not addressed by this policy.
Previous
Next
22% completed
Your Progress
Jump directly to question

Impact on Accuracy

Condition	Correct	Wrong	No Convergence
NOHIGH	76 (84.4%)	4 (4.4%)	10 (11.1%)
TOP05	74 (82.2%)	9 (10 %)	7 (7.8 %)
TOP10	81 (90.0%)	3 (3.3%)	6 (6.7 %)

Performance on 90 policy-question pairs
NOHIGH = No Highlights
TOP 05 = 5 Highlighted Paragraphs
TOP 10 = 10 Highlighted Paragraphs

Note: on average a policy has a little over 40 paragraphs
Suggests possible improvements in both accuracy and productivity

Next Steps

- **Adaptive Crowdsourcing:**
  - Dynamically adjust number of crowdworkers
  - Distinguish between crowdworkers & optimize the allocation of crowdworkers to tasks
- **More organic annotation process**
  - Annotate finer grain issues
- **Public crowdsourcing site: Q1 2016**

Simple Browser Plug-In

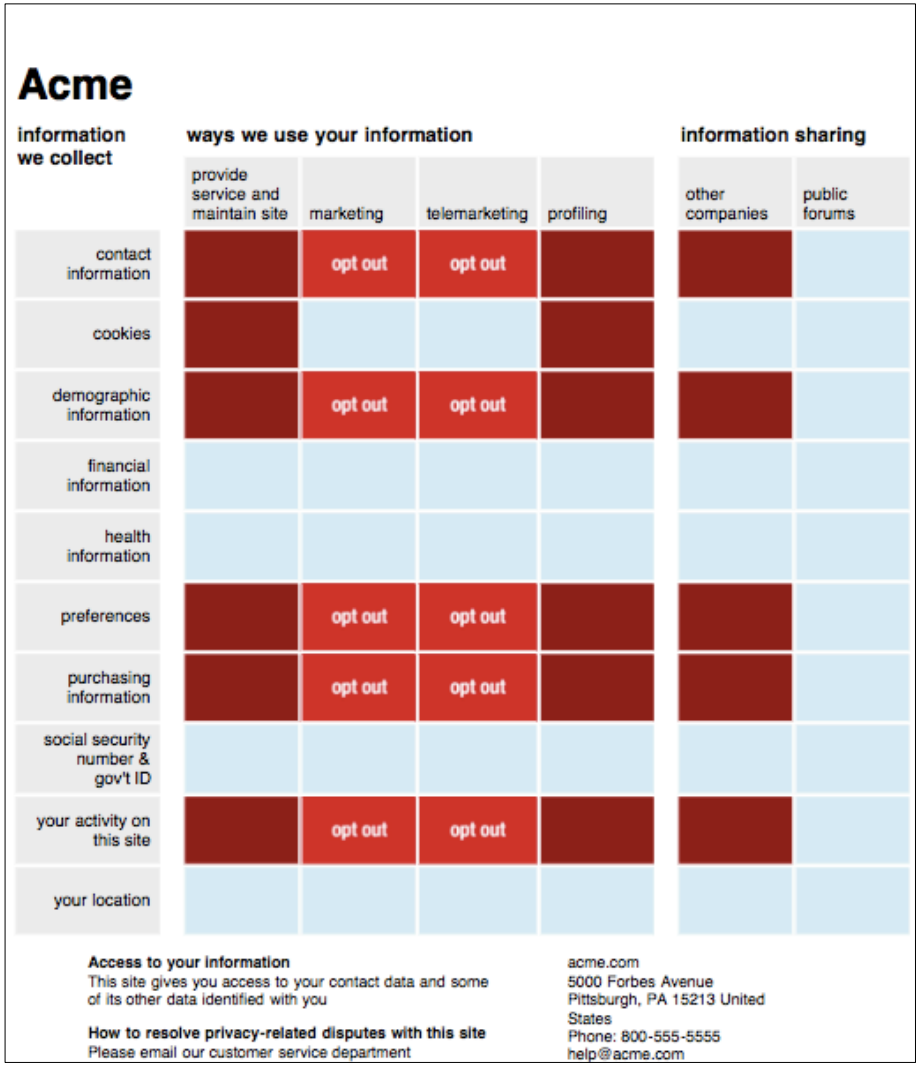

Example of Privacy Nutrition Label

P. G. Kelley, L. Cesca, J. Bresee & L. F. Cranor *Standardizing privacy notices: an online study of the nutrition label approach* CHI '10, ACM 2010.

Informed by Mental Models

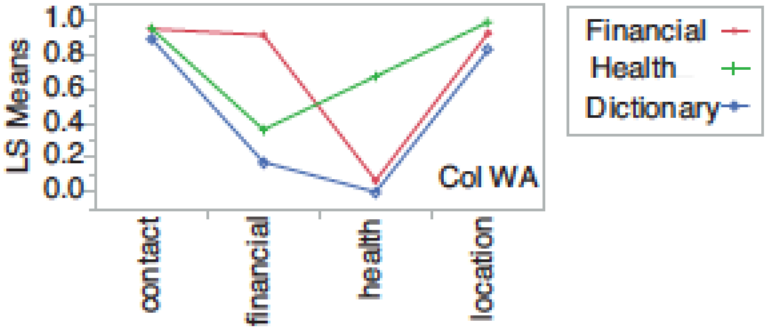

(a) Collection (with account)

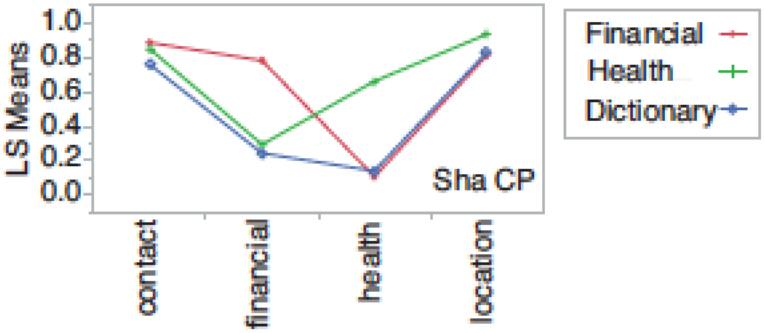

(b) Sharing (core purpose)

Key Insight: Highlight What Users Do Not Expect

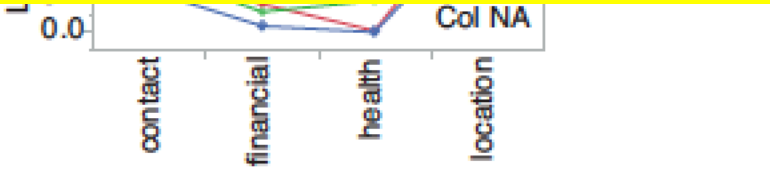

(c) Collection (no account)

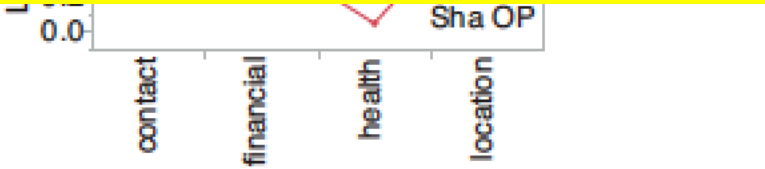

(d) Sharing (other purpose)

User data collection & sharing expectations for 3 categories of websites: financial, health and dictionary

But Simple UI's Many Not Be Enough

….How do we motivate users to pay attention?

Helping Users Manage Privacy Settings: Explosion of Settings

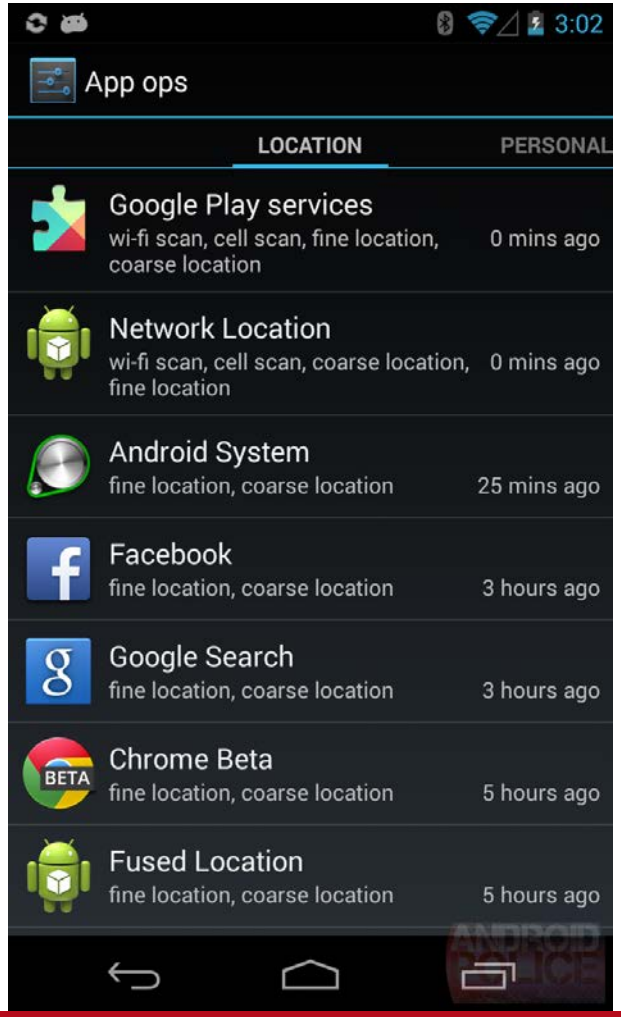

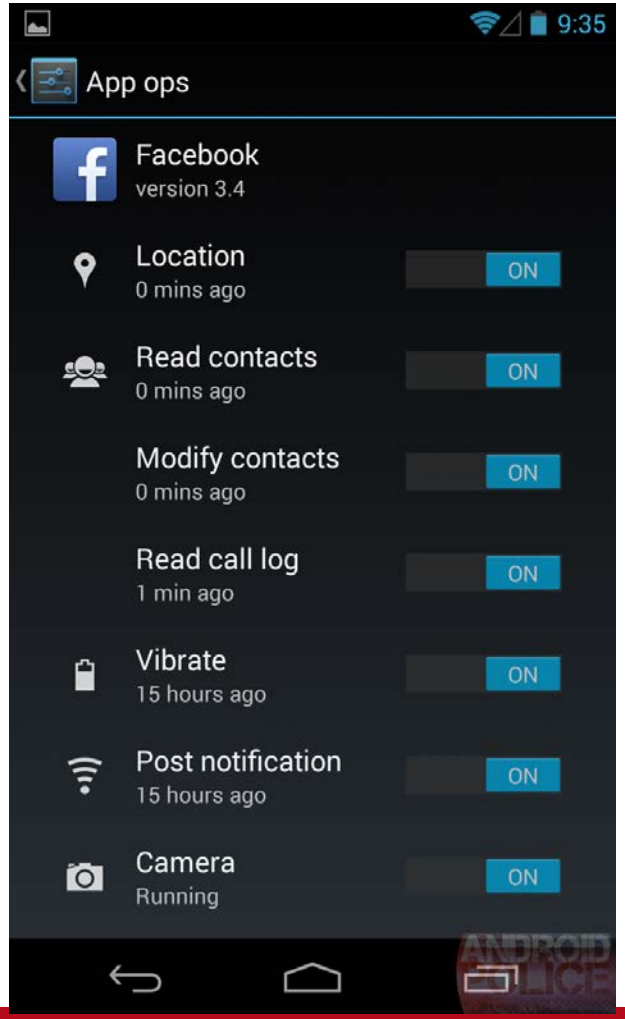

Another Quick Question

- How many of you know what mobile apps are currently running on their smartphones and what information they collect?

People's Response When They Find Out…

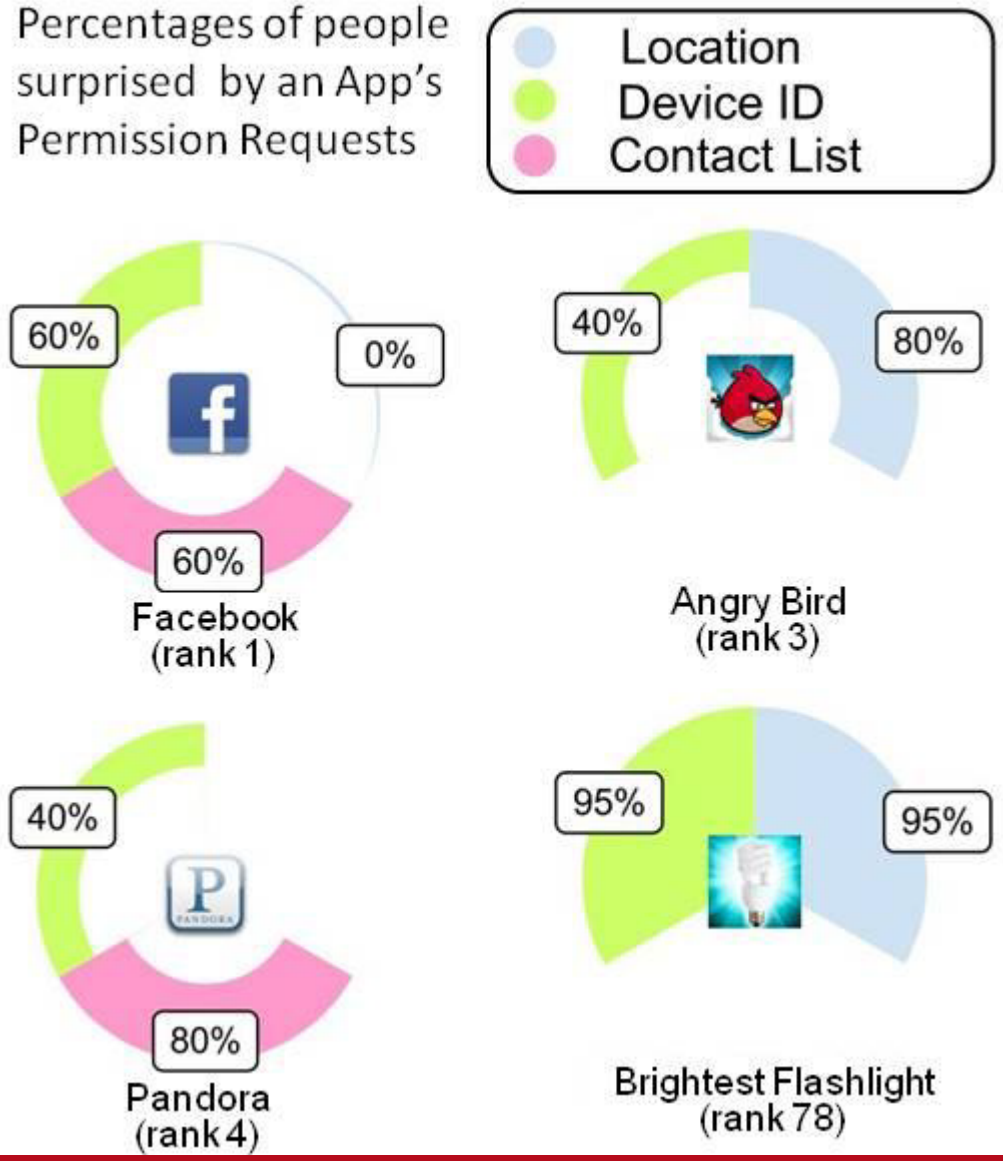

J. Lin, S. Amini, J. Hong, N. Sadeh, J. Lindqvist, J. Zhang, "Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy through Crowdsourcing", Proc. of the 14th ACM International Conference on Ubiquitous Computing, Pittsburgh, USA, Sept. 2012

Helping Users Manage Privacy Settings

- …Beyond understanding privacy policies…
- Do permission managers help users (e.g. iOS, Android App Ops)?
- **Could nudges help increase user awareness and motivate users to take a closer look at their settings?**

Field Study

- 22-day study with 23 participants using their regular Android phone
- **Week 1:** baseline
- **Week 2:** App Ops permission manager
- **Final 8 days:** App Ops + one daily nudge focused on one permission
- Collected detailed logs of all permission changes + pre- and post-surveys

Privacy Nudge

Detailed Report

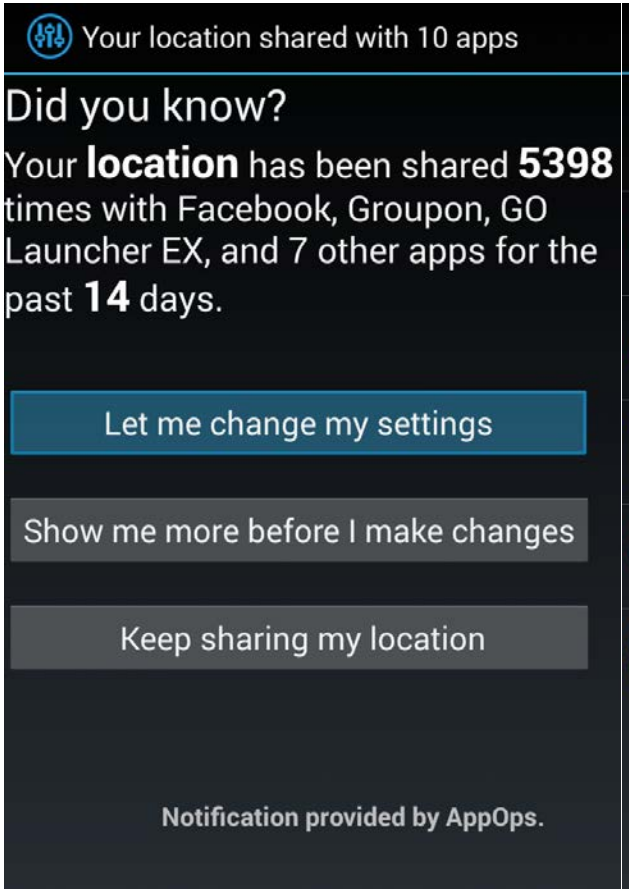

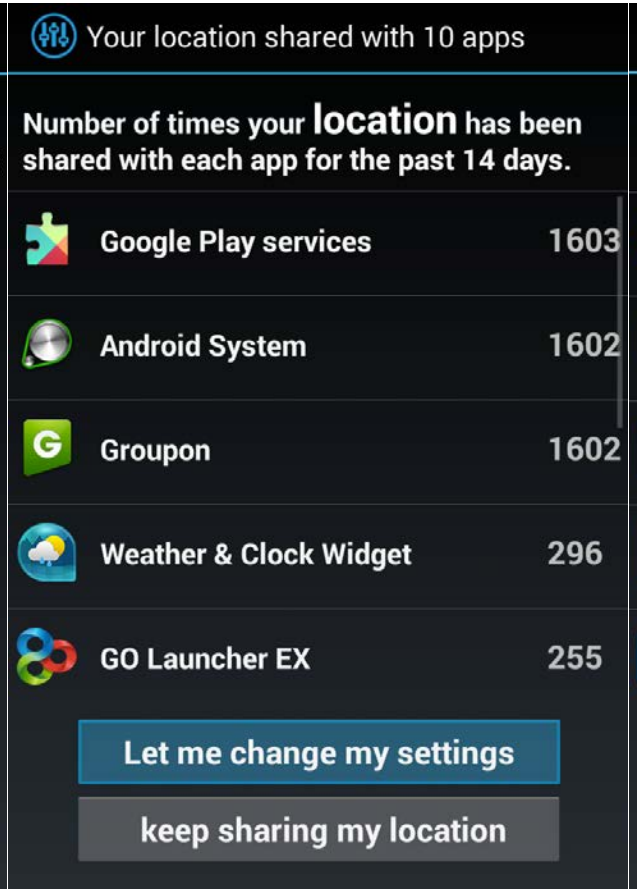

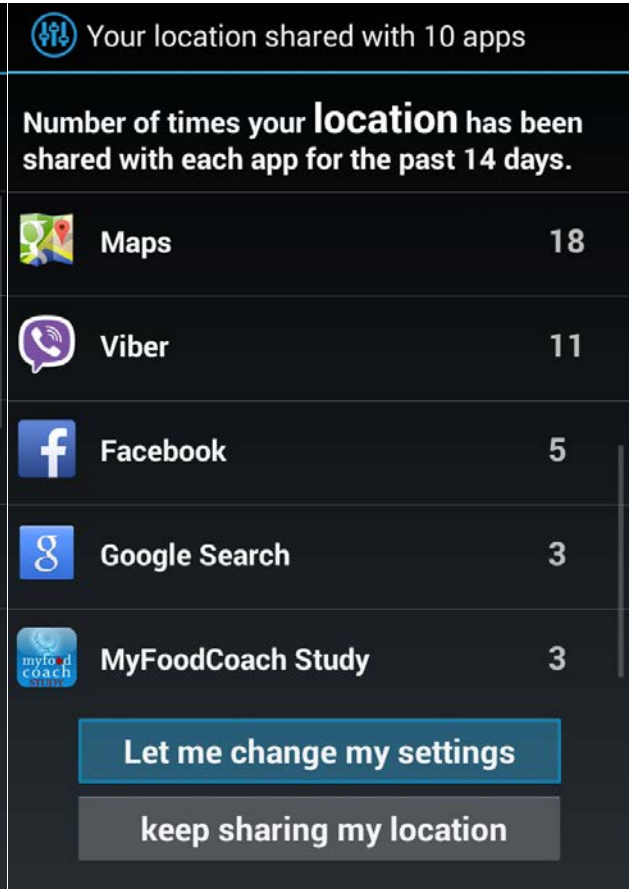

Implementation of Study App

- **Developed & installed an Android app** that was used:
  - To **launch AppOps**
  - **Collect detailed AppOps logs**
    - What permissions apps are allowed to access
      - Including changes to settings made by user, when and through which interface (e.g. from the nudge or directly via AppOps)
    - For each app-permission pair:
      - Last time the app tried to access the permission
      - Each app request & whether it was granted/denied
      - Whether the app is currently using the permission and for how long (e.g. camera, recording audio)

Demographics & Additional Details

- **23 participants** (65% female; ages 18–44, median=23)
  - 21 owned Samsung devices and 2 owned an HTC One.
- On average, 89 apps installed (SD=22), including services and pre-installed apps.
- **21 (91%) reported never using AppOps before**
  - 1 had used AppOps, and 1 was unsure.
- **Phase 1 users could not access App Ops**
    - *We checked that no other App Ops launcher was installed on their phones and our App did not allow them during that week to access App Ops*.

Permission Managers Are Not Enough

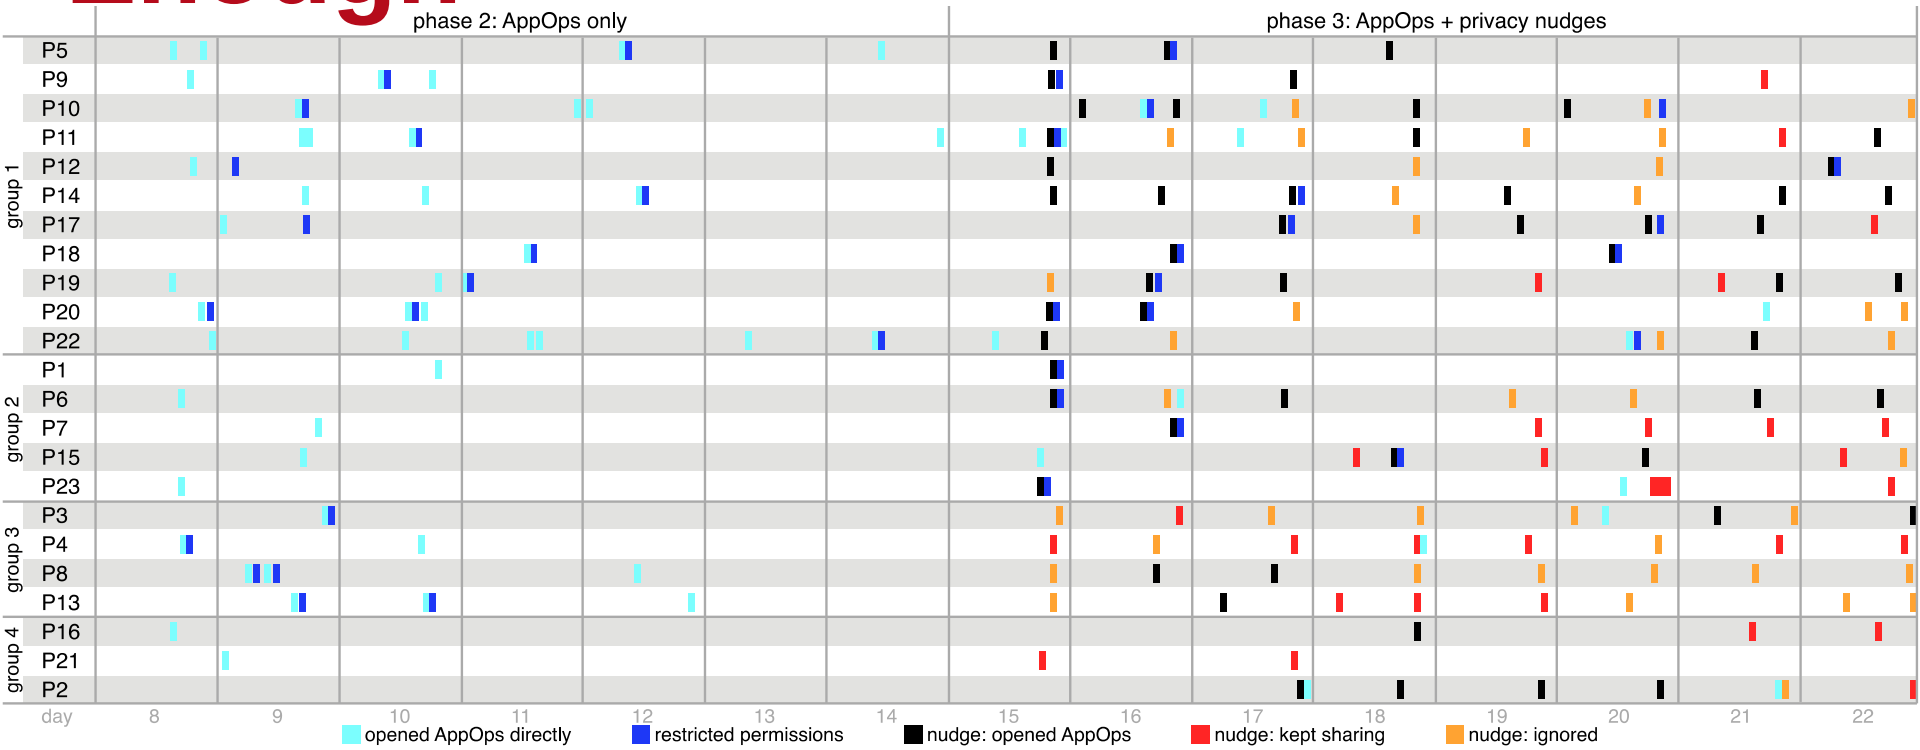

Week 2: Permission Manager Only **Week 3: Daily Nudges**

- Nudges can make a big difference

Permission Manager w/o Privacy Nudges

- In phase 2, participants reviewed their app permissions **51 times**, restricted **76 distinct apps** from accessing a total of **272 permissions**
- Only one interaction where a user opened access to one permission.

Reviewing of App Permissions During Phase 2

- **22 participants (95.6%) reviewed their app permissions at least once**.
  - 12 reviewed their app permissions multiple times.
- One did not review his permissions in phase 2.

Adjusting App Permissions

- 15 (**65%**) participants restricted **272 app-permission pairs from 76 distinct apps**, including both participant-installed and pre-installed apps.
- Participants restricted apps' access to:
  - **Location:** 74 (27%)
  - **Contacts**: 57 (21%)
  - **Calendar**: 10 (4%)
  - **Call logs**: 9 (3%).
  - Others included: Camera: 42 (9%), SMS: 21 (8%), Post notification: 19 (7%), Recording audio: 15 (6%).
- Only one participant opened back a permission to the Weather Channel app - to send notifications.

Why did participants restrict apps' access to permissions?

- Participants **restricted unused apps, especially pre-installed apps**.
  - P10 stated: "I also blocked [a] bunch of AT&T bloatware from accessing any information. I don't use them anyways."
- Participants restricted permissions required for unused functionality.
  - P13 restricted iHeartRadio access to location, explaining: "I know what stations I want to listen to no matter where I am so I turn off the location."

Why did participants restrict apps' access to permissions? (continued)

- Participants **restricted apps when the purpose to access their personal information was unclear**.
  - P4 stated: "[I turned it off] because I can't think of a reason why Inkpad needs my location."

Adding Privacy Nudges – Final 8 days

- Do nudges further change user behavior and how do they feel about them?

Effectiveness of Privacy Nudges – Final 8 days

- In phase 3, participants reviewed their app permission **69 times**, restricted **47 distinct** apps from accessing **122 permissions**, and permitted six apps access to six permissions.

- ….**this is after a week with access to App Ops.**

Reviewing App Permissions

- Participants could review their app permissions either by
  1. **Opening AppOps directly** (same as in phase 2)
  2. **Opening AppOps in response to a nudge**.

Reviewing App Permissions

- 22 participants (95.6%) reviewed their app permissions at least once in phase 3.
  - **21 participants reviewed their apps' permisions <u>in response to nudges</u>:**
    - 53 times (**78% of the time) in response to a nudge**
    - 15 times (22%) by directly opening AppOps.
  - **1 participant** reviewed her apps' permissions only once and **only by directly opening AppOps**
- **The privacy nudges were the primary trigger for participants to review their app permissions.**

Adjusting App Permissions

- Participants restricted 122 permissions - breakdown:
  - **Location: 30 (25%)**
  - **Contacts: 25 (20%)**
  - **Calendar: 8 (7%)**
  - **Call logs: 6 (5%).**
  - Other restricted permissions included:
    - Post notification: 10 (8%), SMS: 9 (7%), camera: 7 (6%), record audio: 7 (6%).
- Only three participants made **permissive adjustments due to loss of app functionality**.
  - In the interview, P10 noted that he restricted and later permitted Facebook's access to the clipboard, because he was unable to copy&paste in Facebook.

Concluding Remarks - I

- "Notice and Choice" is the de facto approach to privacy on the Web
- Even on the fixed Web, this approach does not work
- On smartphones and with the emerging Internet of Things, this **framework (in its current form) simply does not scale**

Concluding Remarks - II

- **Usability research** is key to understanding **what users can realistically be expected to do** and can inform the design of more realistic implementations of Notice and Choice
- **Artificial Intelligence**, which is often blamed for many of the privacy risks we face in the Internet of Things, may also hold part of the solution to **scaling Notice and Choice**
  - Machine Learning & NLP to **annotate privacy policies**
  - Personalized Privacy Assistants for **personalized summaries of policies**, **learning our privacy preferences** and **nudging** us to examine our privacy settings

Apps Snoop on Your Location Way More Than You Think

Wired - Mar 25, 2015

But Professor **Norman Sadeh**, a member of the research team that conducted the research, says the volume of location-harvesting isn't the ...

What do your mobile apps tell third parties?

Futurity: Research News - Mar 26, 2015

Carnegie Mellon study: your apps are tracking you more than you ...

ConsumerAffairs - Mar 24, 2015

How Often Does an App Share Your Location? You May Be Surprised

Credit.com News (blog) - Mar 24, 2015

You'll Be Freaked Out to Learn How Often Your Apps Share Your ...

Money - Mar 25, 2015

Where Were You 3 Minutes Ago? Your Apps Know

Highly Cited - Wall Street Journal (blog) - Mar 23, 2015

ConsumerAff...

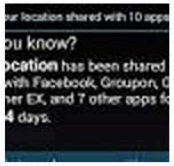

Wall Street J...

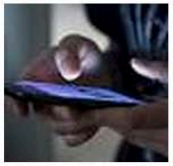

Credit.com N...

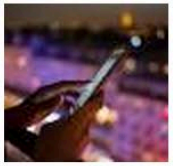

Money

BGR

Explore in depth (37 more articles)

Does Your Cellphone Know Too Much?

90.5 WESA - Apr 6, 2015

... private information, they tried to limit future sharing, according to **Norman Sadeh**, a Carnegie Mellon University computer science professor.

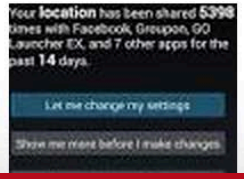

Your location has been shared 5398 times …. 'Are you kidding me?'

Bob Sullivan.net - Mar 27, 2015

“The vast majority of people have no clue about what's going on,” said **Norman Sadeh**, a professor in the School of Computer Science's

Acknowledgements: Work funded by the National Science Foundation, Google and Samsung

The **Usable Privacy Policy Project** involves a collaboration with a number of individuals. See **usableprivacy.org** for additional details

Q&A