Artificial Intelligence, the Internet of Things and Privacy: ***Are We Doomed?***

Norman Sadeh

Carnegie Mellon University

www.normsadeh.org

usableprivacy.org privacyassistant.org

explore.usableprivacy.org

What is AI?

- A branch of Computer Science encompassing a vast collection of technologies aimed at **emulating, supplementing, and eventually outperforming human intelligence**
  - From knowledge representation and reasoning to computer vision, robot planning, intelligent assistants, etc.
  - From ”symbolic” (and logic-based) frameworks to neural networks
- Includes **machine learning** and data mining

AI Bragging Rights (Short Selection)

Today AI is Everywhere

- Translation
- Scene recognition
- Facial expression recognition – incl. recognizing emotions
- Autonomous driving, autonomous flying of drones, etc
- Speech recognition
- Automated trading
- Discovering new uses for existing drugs
- Analyzing DNA to detect genomic conditions
- Detecting crop diseases and predicting crop yield
- Help users book tables at restaurant and movie tickets
- …and much more…

AI is Powered by Data

- AI and in particular machine learning is **powered by data**
- The more data you collect, the more patterns you might be able to recognize
  - Richer sets of predictions
  - Prediction accuracy

IoT and AI

- The Internet of Things is about embedding processing, communication, sensing and actuation functionality in everyday objects
- **Every object around you can now be a source of data and can also leverage inferences derived from data**
  - Everything can talk to everything else

Should We Be Concerned About our Privacy?

- Increasingly **diverse, complex and opaque** data collection & use processes
  - Secondary use of data
  - Data being recombined
  - Data mining: difficult to control and to explain
- **91%** of people report feeling they **have lost control over their information**

  Pew Survey 2014 http://www.pewinternet.org/2014/11/12/public-privacy-perceptions/

…And What is Privacy Anyway?

- Moral right of individuals to **be left alone, free from surveillance or interference from other individuals or organizations**, including state
  - There are obviously **conflicting considerations**
    - e.g. security and safety, personalization, productivity

Classes of Privacy

- **Information Privacy:** collection and handling of information
- **Bodily Privacy:** includes drug testing, strip search, abortion, adoption, etc.
- **Territorial Privacy:** limits on the ability to intrude into another individual's environment
- **Communications Privacy:** includes postal mail, email, telephone, etc.

Information Privacy

- The claim that ***certain information should not be collected by government or businesses*** – *or possibly* ***only under special circumstances*** and subject to various rules
  - **individuals have some control over the use of information collected about them**

Legal & Regulatory Landscape

- A number of privacy laws around the world:
  - Hong Kong Personal Data Ordinance
  - EU General Data Protection Directive
  - US: patchwork of sectoral laws
- All these laws share some **commonalities: They set minimum requirements to:**
  - **Inform** users about data collection and use practices
  - Provide users with some type of **choice**

Challenges

- **Tension between Privacy and User Burden**
  - Privacy as a **Secondary Task**
  - Lack of **awareness**
  - **Cognitive and Behavioral Biases**
  - **Explosion** in Number of Privacy **Decisions**
- Increasingly **diverse, complex and opaque data collection and use practices**

AI to the Rescue

- **Regulations like GDPR help** by requiring more **explicit disclosure and consent** processes (among other things) and by introducing **steep penalties**
- The **sheer scale** at which data is being collected and used in the age of AI and IoT **requires that AI also be used to help users, companies and regulators**

People Never Read Privacy Policies

- Could we get computers to do it for them?
  - And in the process also help
    - Data collectors, processors and regulators verify that policies are compliant

Wilson, S., Schaub, F., Ramanath, R., Sadeh, N., Liu, F., Smith, N., and Liu, F. Crowdsourcing Annotations for Websites Privacy Policies: Can It Really Work? WWW Conference, May 2016

We Spent a Lot of Time Collecting Policy Annotations

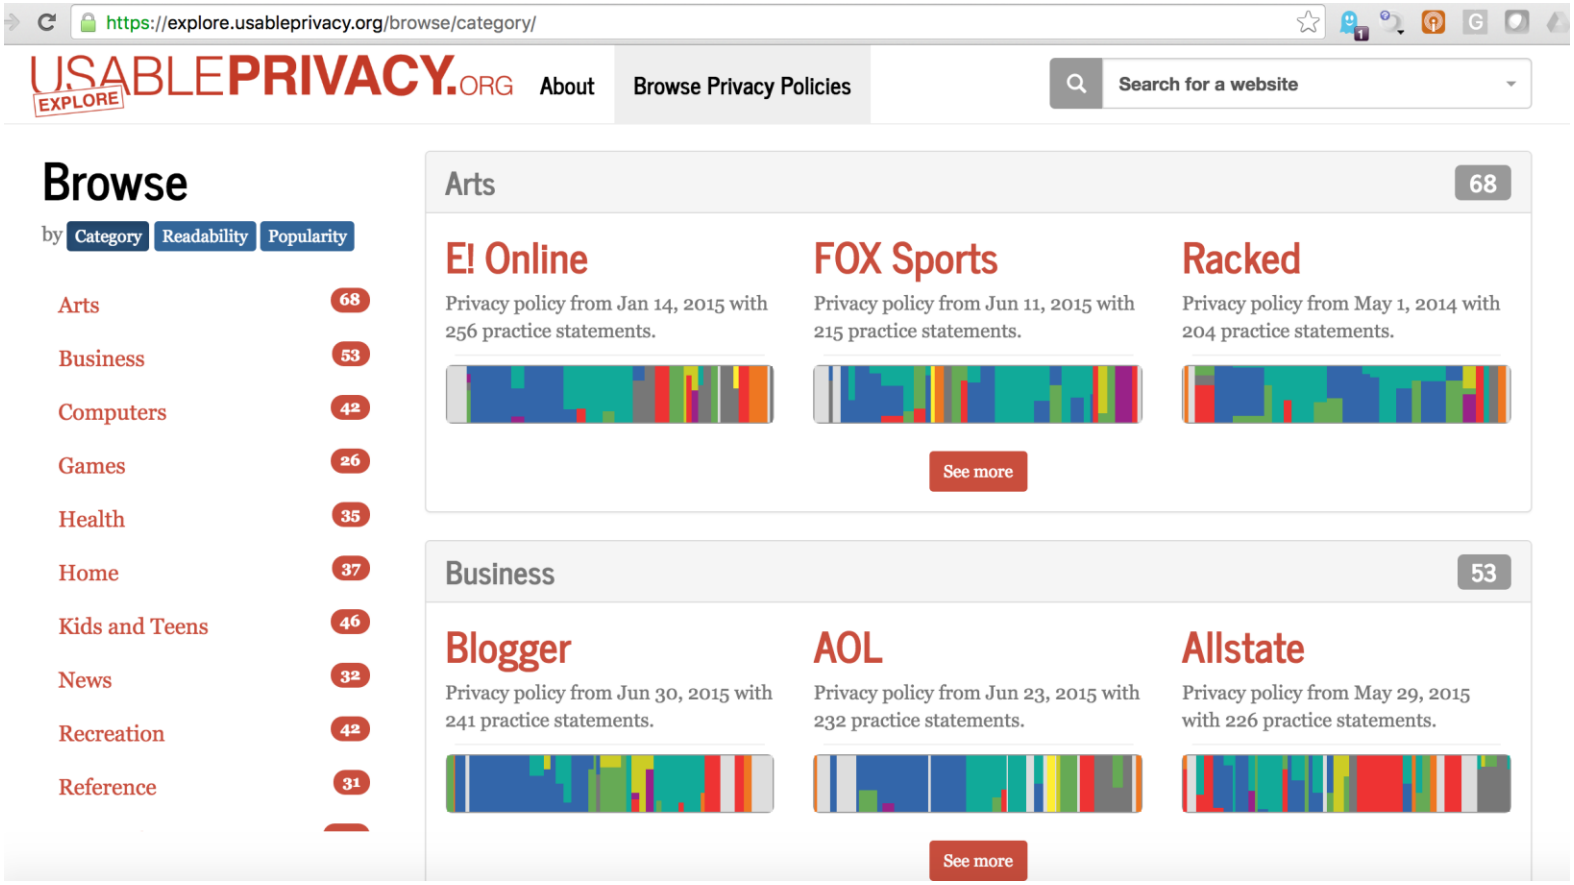

S. Wilson, F. Schaub, A. Dara, F. Liu, S. Cherivirala, P.G. Leon, M.S. Andersen, S. Zimmeck, K. Sathyendra, N.C. Russell, T.B. Norton, E. Hovy, J.R. Reidenberg, N. Sadeh, "The Creation and Analysis of a Website Privacy Policy Corpus", ACL '16: Annual Meeting of the Association for Computational Linguistics, Aug 2016

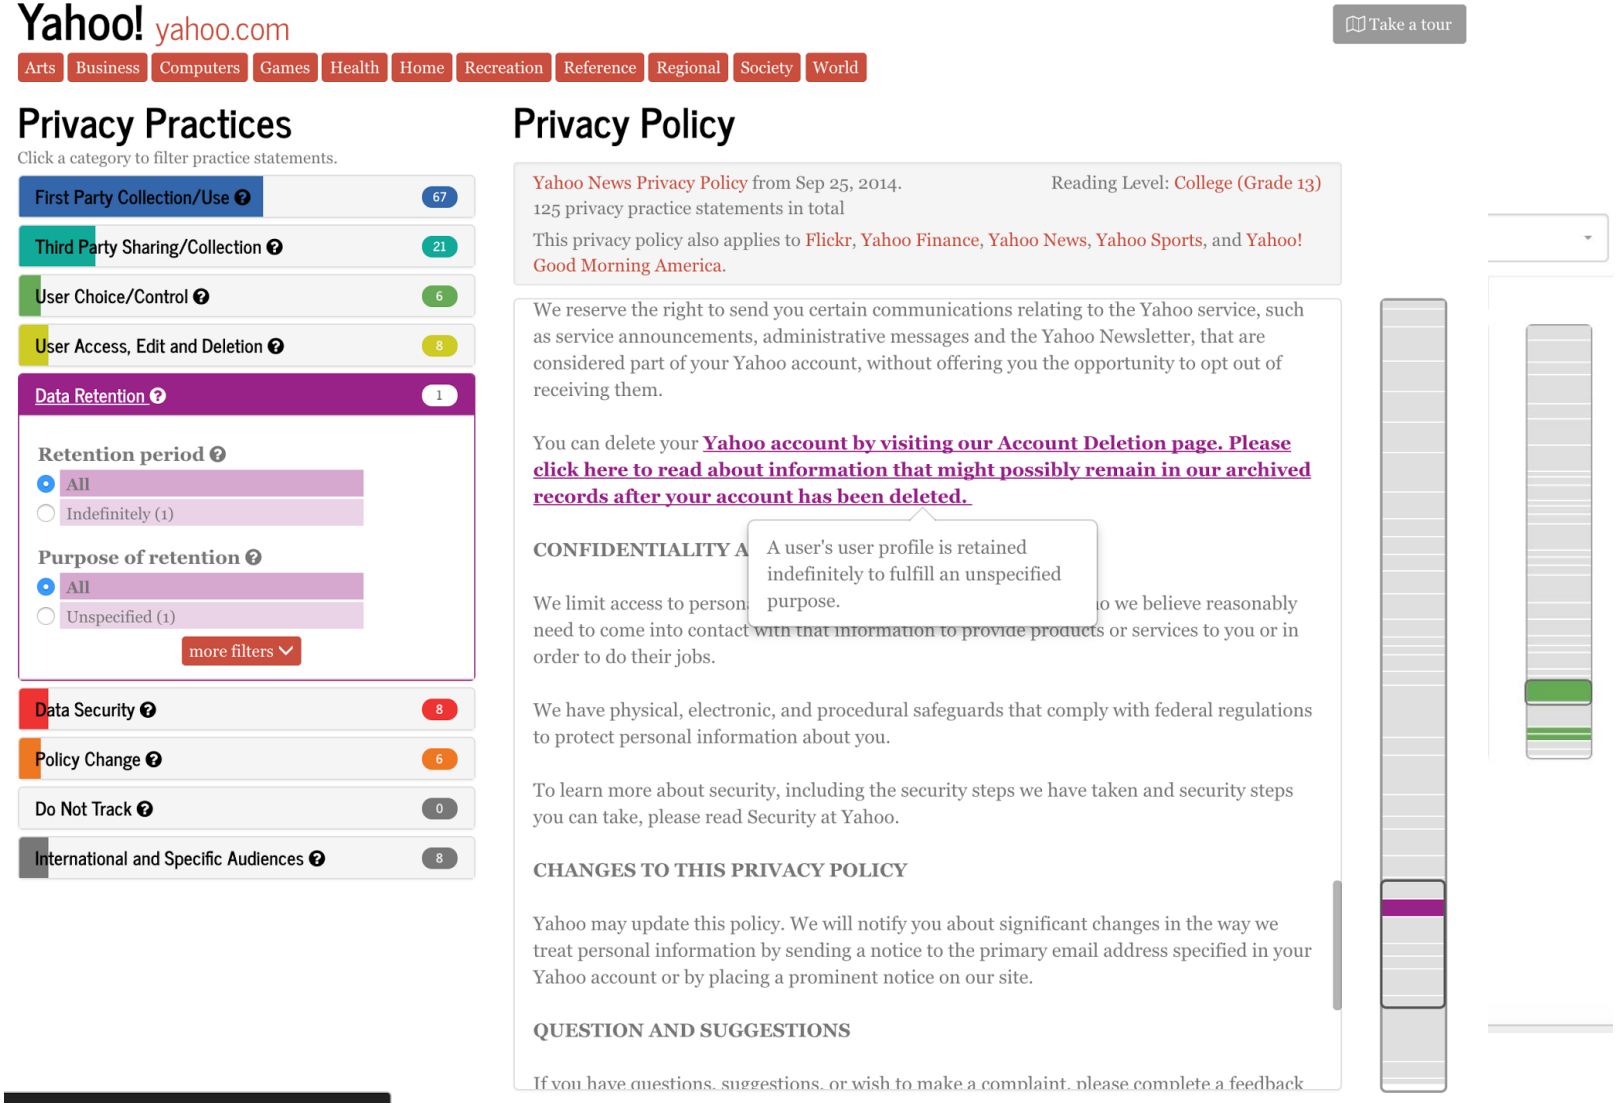
Yahoo! yahoo.com
Take a tour
Arts Business Computers Games Health Home Recreation Reference Regional Society World
Privacy Practices
Click a category to filter practice statements.
First Party Collection/Use 67
Third Party Sharing/Collection 21
User Choice/Control 6
User Access, Edit and Deletion 8
Data Retention 1
Retention period
All
Indefinitely (1)
Purpose of retention
All
Unspecified (1)
more filters
Data Security 8
Policy Change 6
Do Not Track 0
International and Specific Audiences 8
Privacy Policy
Yahoo News Privacy Policy from Sep 25, 2014.
Reading Level: College (Grade 13)
125 privacy practice statements in total
This privacy policy also applies to Flickr, Yahoo Finance, Yahoo News, Yahoo Sports, and Yahoo! Good Morning America.
We reserve the right to send you certain communications relating to the Yahoo service, such as service announcements, administrative messages and the Yahoo Newsletter, that are considered part of your Yahoo account, without offering you the opportunity to opt out of receiving them.
You can delete your Yahoo account by visiting our Account Deletion page. Please click here to read about information that might possibly remain in our archived records after your account has been deleted.
A user's user profile is retained indefinitely to fulfill an unspecified purpose.
CONFIDENTIALITY A
We limit access to person
o we believe reasonably need to come into contact with that information to provide products or services to you or in order to do their jobs.
We have physical, electronic, and procedural safeguards that comply with federal regulations to protect personal information about you.
To learn more about security, including the security steps we have taken and security steps you can take, please read Security at Yahoo.
CHANGES TO THIS PRIVACY POLICY
Yahoo may update this policy. We will notify you about significant changes in the way we treat personal information by sending a notice to the primary email address specified in your Yahoo account or by placing a prominent notice on our site.
QUESTION AND SUGGESTIONS
If you have questions, suggestions, or wish to make a complaint, please complete a feedback

Automatic Segment Annotation

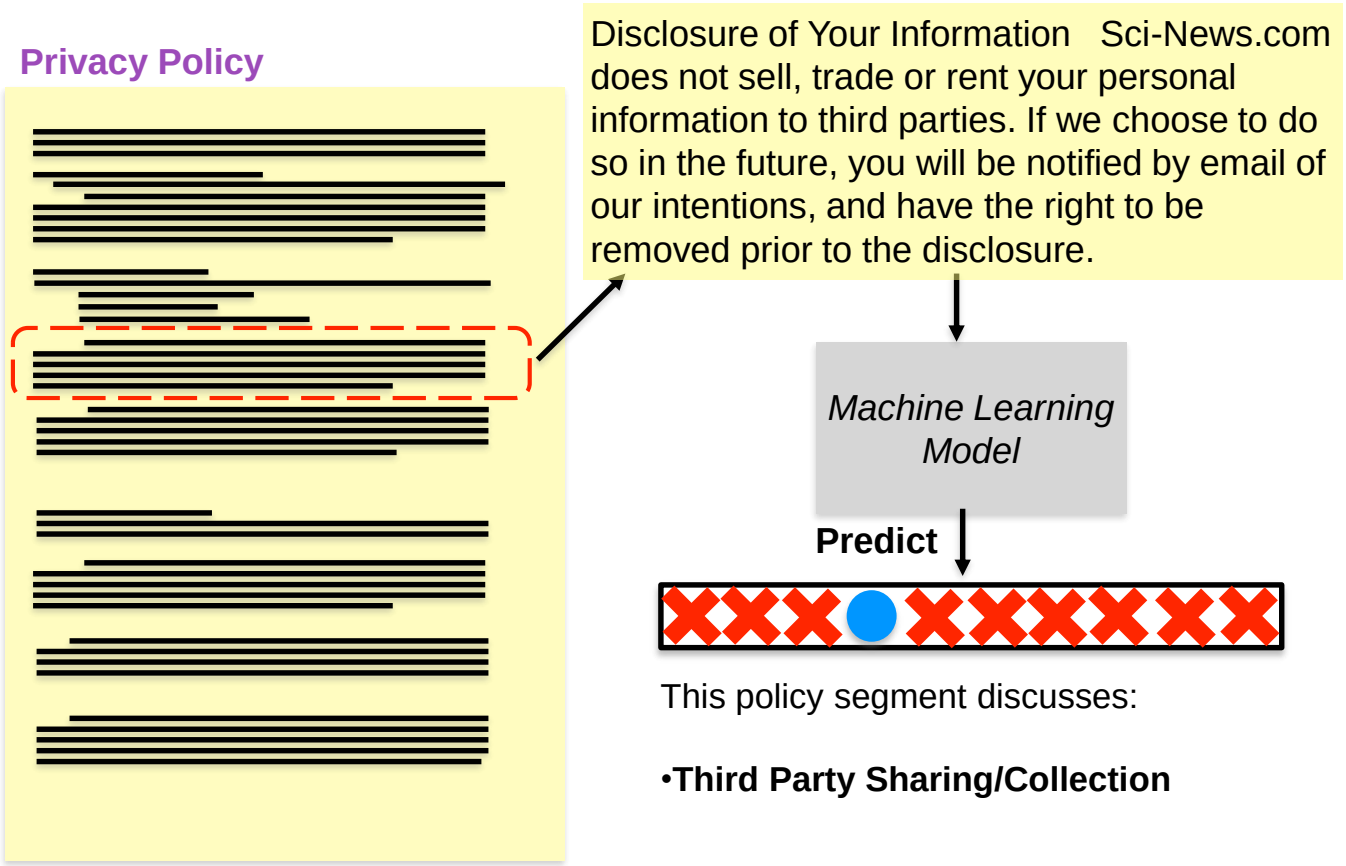

F. Liu, S. Wilson, F. Schaub, N. Sadeh.. *Analyzing Vocabulary Intersections of Expert Annotations and Topic Models for Data Practices in Privacy Policies* AAAI Fall Symposium on Privacy and Language Technologies. 2016.

Automatic User Choice Instance Extraction

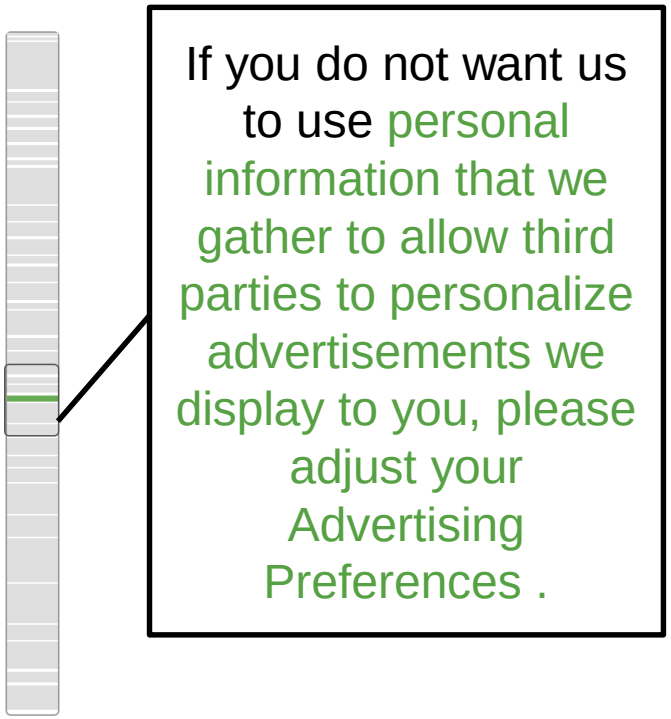

- Users choices often buried deep in the text of long policies
- Decent success at **automatically extracting information** about such "choice instances" from privacy policies?

K.M. Sathyendra, F. Schaub, S. Wilson, N. Sadeh. *Automatic Extraction of Opt-Out Choices from Privacy Policies.* AAAI Fall Symposium on Privacy and Language Technologies. 2016.
K.M. Sathyendra, S. Wilson, F. Schaub, S. Zimmeck, N. Sadeh. *Identifying the Provision of Choices in Privacy Policies, EMNLP Conference, 2017*

Annotated 7,000+ policies

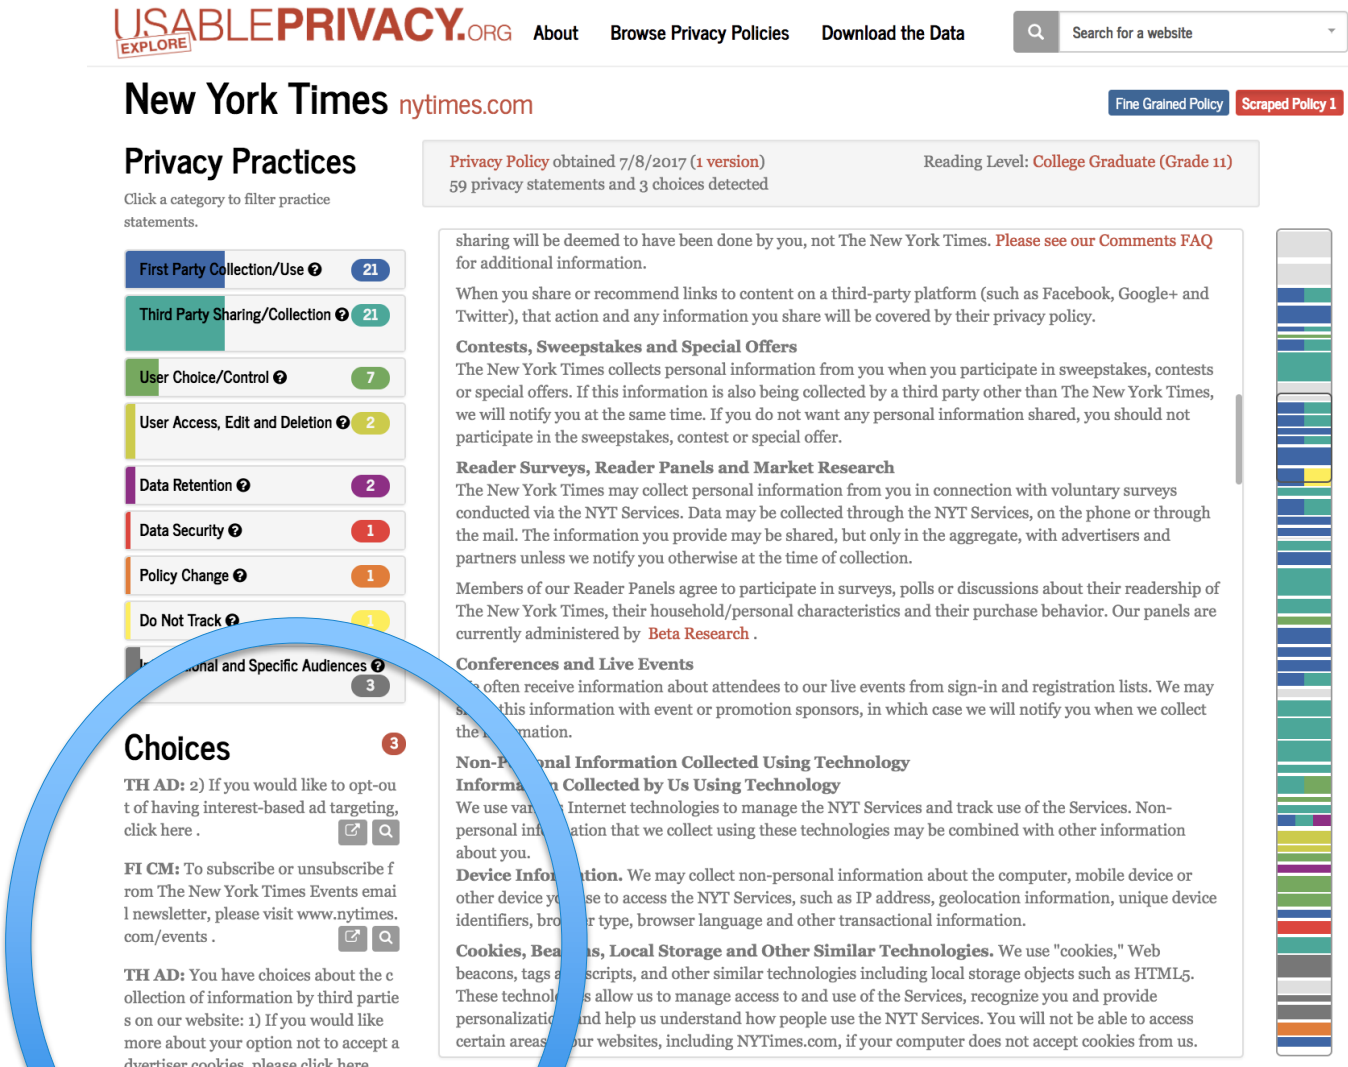

https://explore.usableprivacy.org/

Towards Automated Compliance Analysis

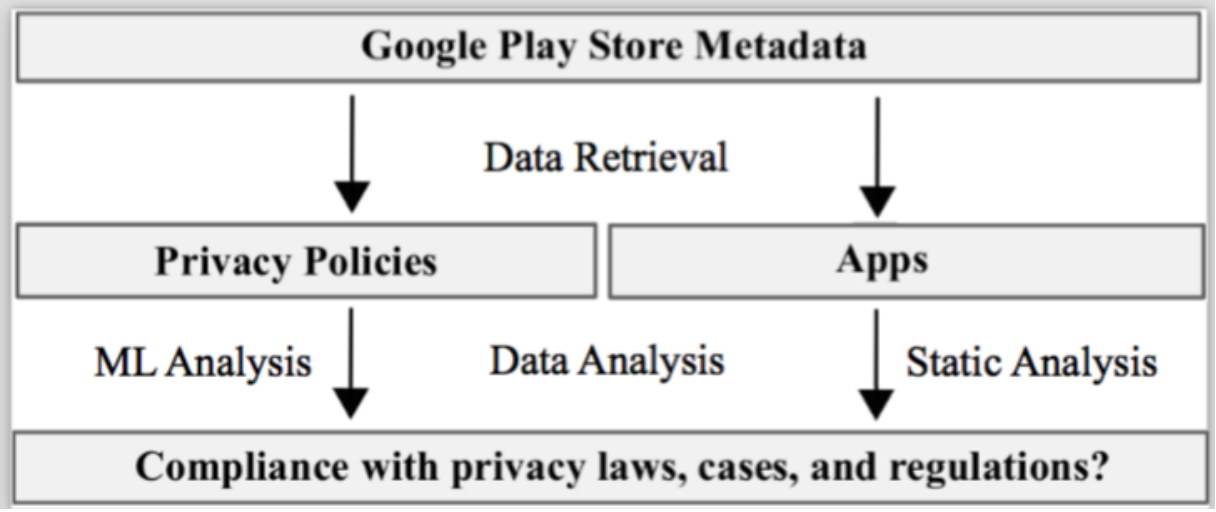

- Training **machine learning classifiers** to extract relevant policy statements
- Compare these statements against:
  - **Regulatory requirements**
  - What the software actually does
    - **Static (and a little bit of dynamic) code analysis**

"Analyzing and Predicting Privacy Law Compliance of Mobile Apps", S. Zimmeck, Z.Wang, L. Zou, B. Liu, F. Schaub, S. Wilson, N. Sadeh, S. Bellovin, J. Reidenberg PLT2016 – longer version at NDSS2017

Formalizing the Problem

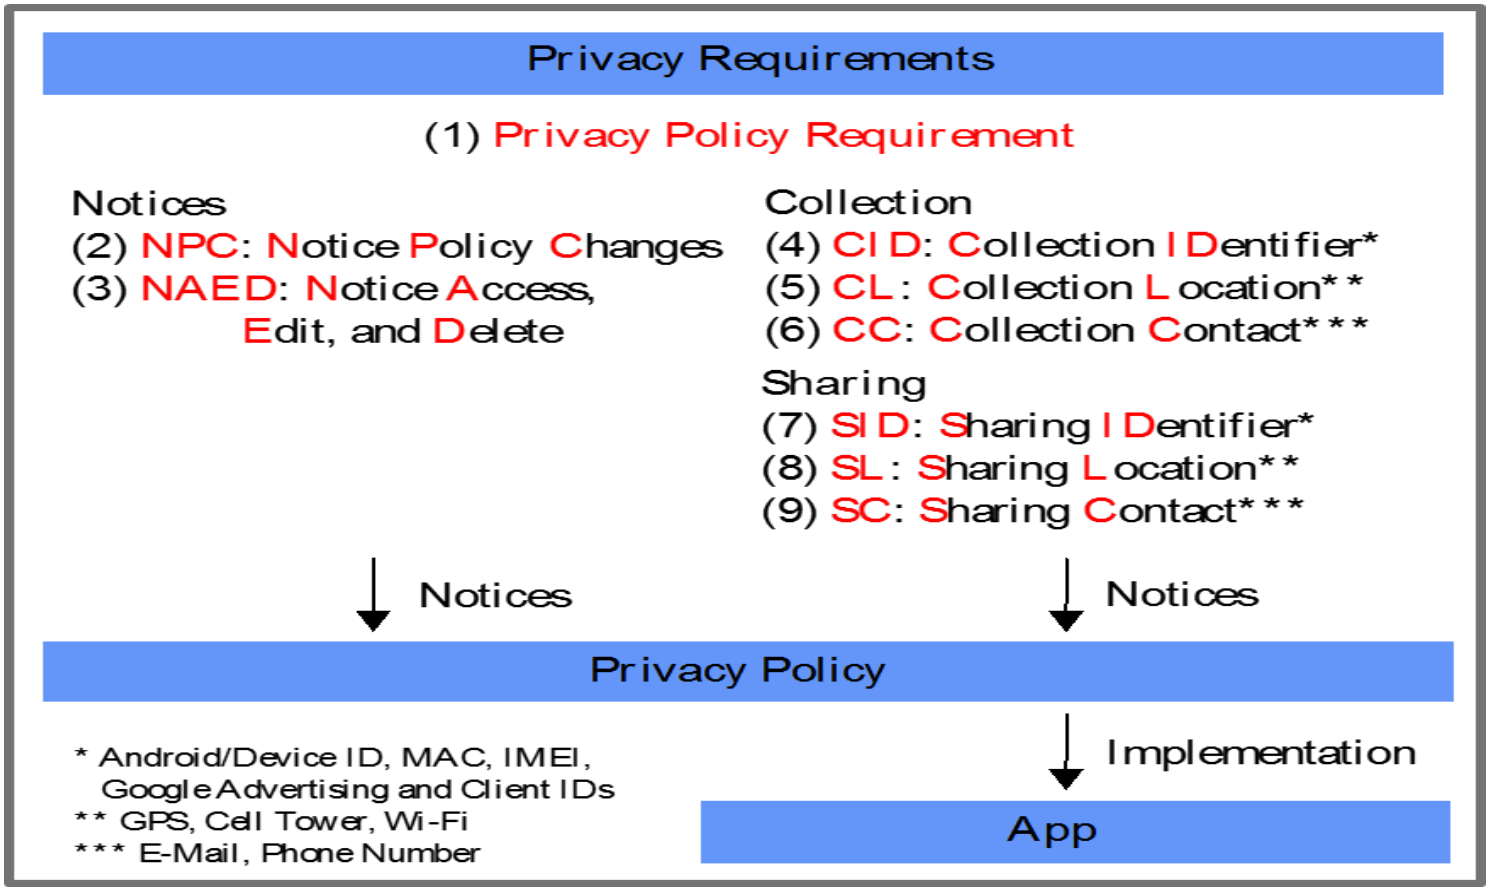

Note: In US, FTC FIPPS mandates notice before collection of PII; COPPA requires policies for apps directed to children; CalOPPA: policy required if PII collected; COPPA requires NAED; CID and CL require disclosure under CalOPPA and COPPA and sharing requires consent; CalOPPA and DOPPA require description of notification process for policy change

People Don’t Take the Time to Review their Privacy Settings

- Could we get computers to help them?

Explosion in Number of Settings

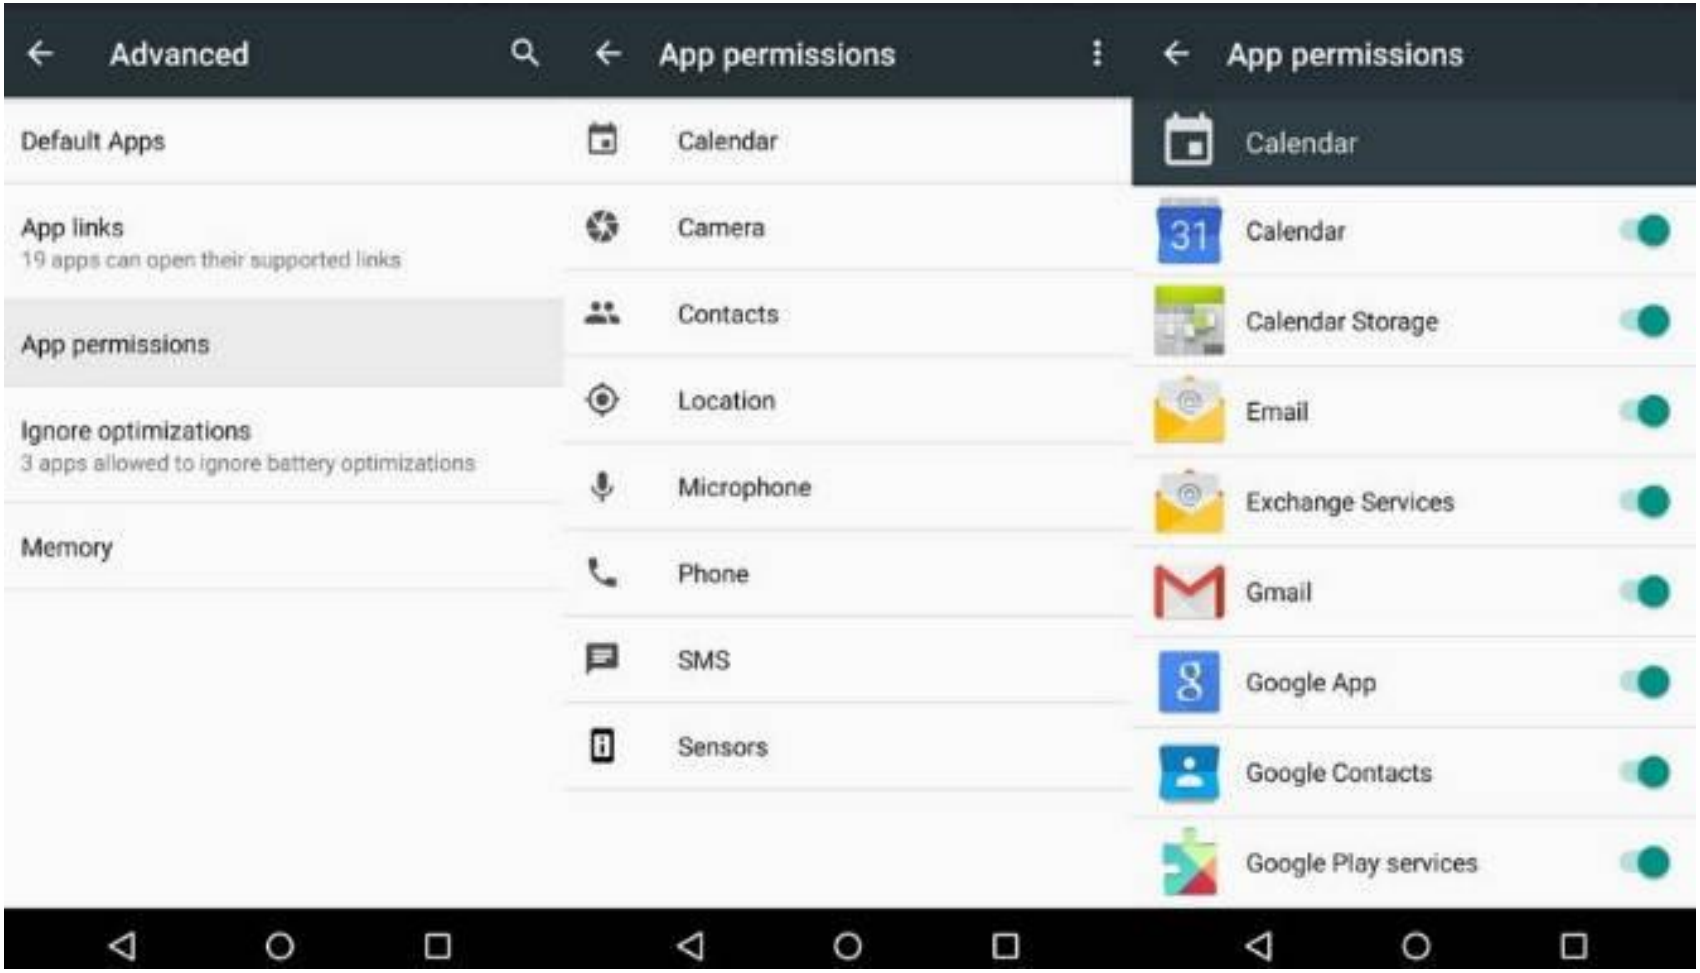

Identifying a User's Privacy Profile

- Using Clustering techniques
- Asking users a small set of questions

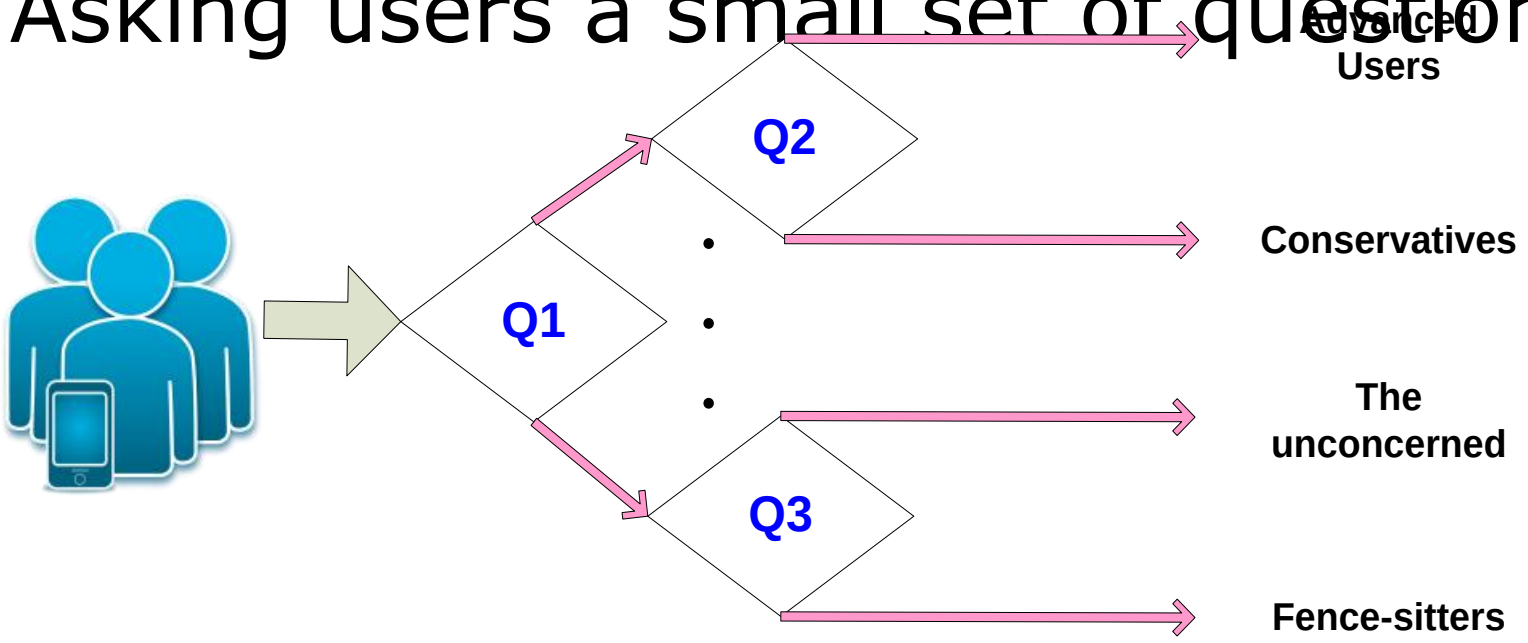

Results with Just 4 Clusters

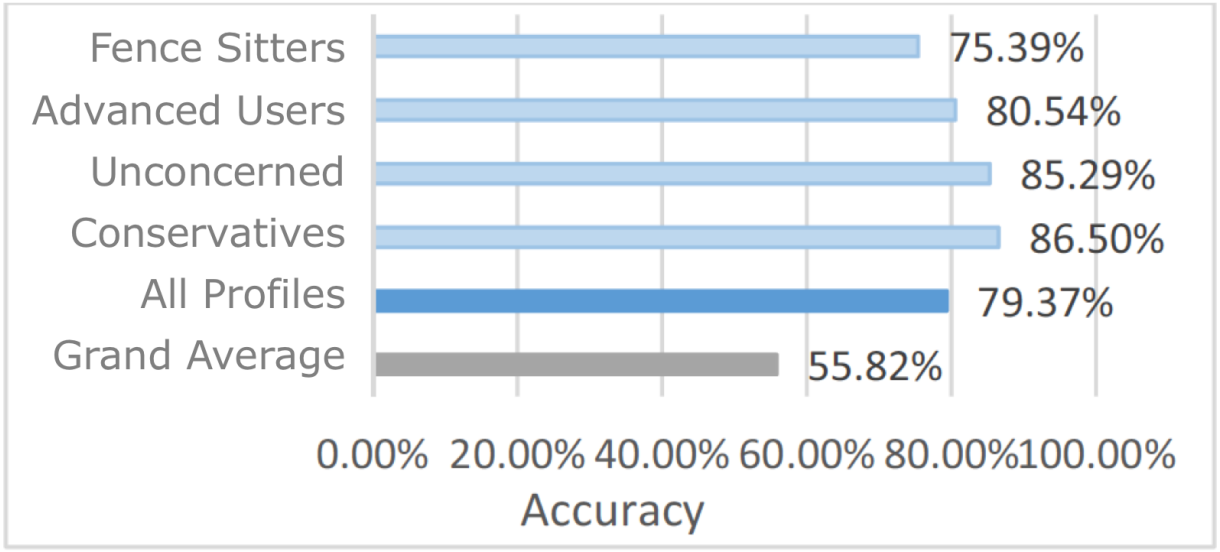

Accuracy:
One size fits all: 55.8%
4 Profiles: 79.4%

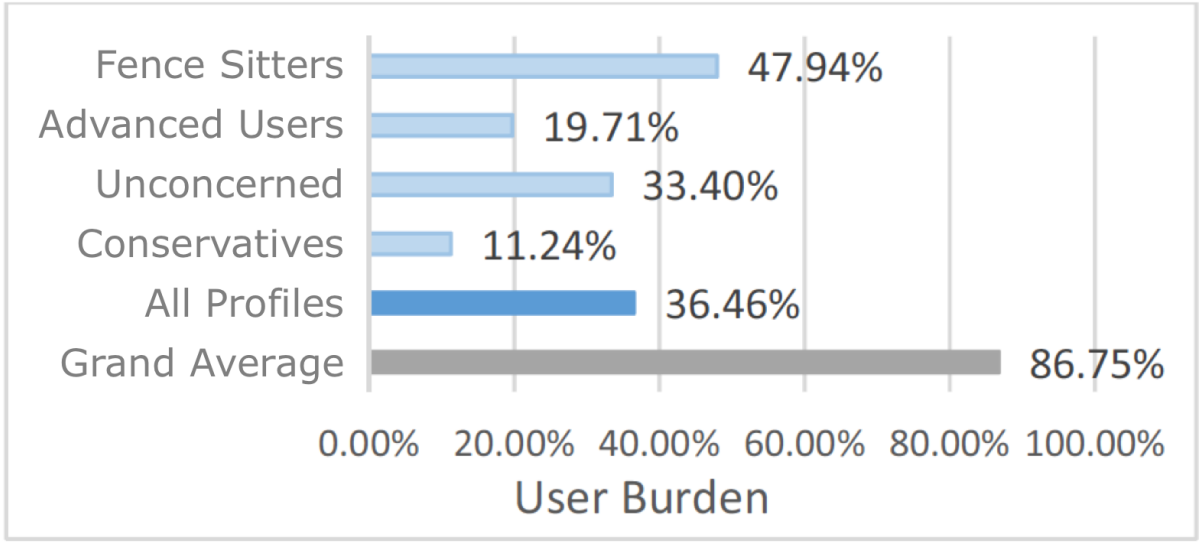

User Burden:
One size fits all: 86.8%
4 Profiles: 36.5%

Now Available on Google Play (rooted Android Phones 5 and up)

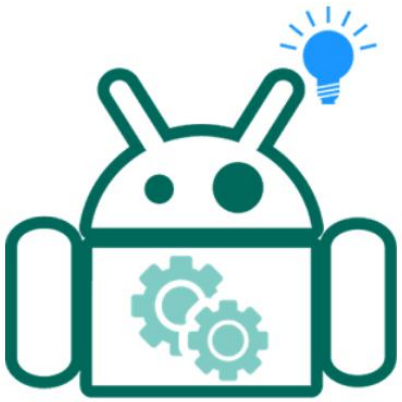

[ROOT] Privacy Assistant

Mobile Commerce Lab @ Carnegie Mellon University **Tools** ★★★★☆ 4

E **Everyone**

Add to Wishlist **Install**

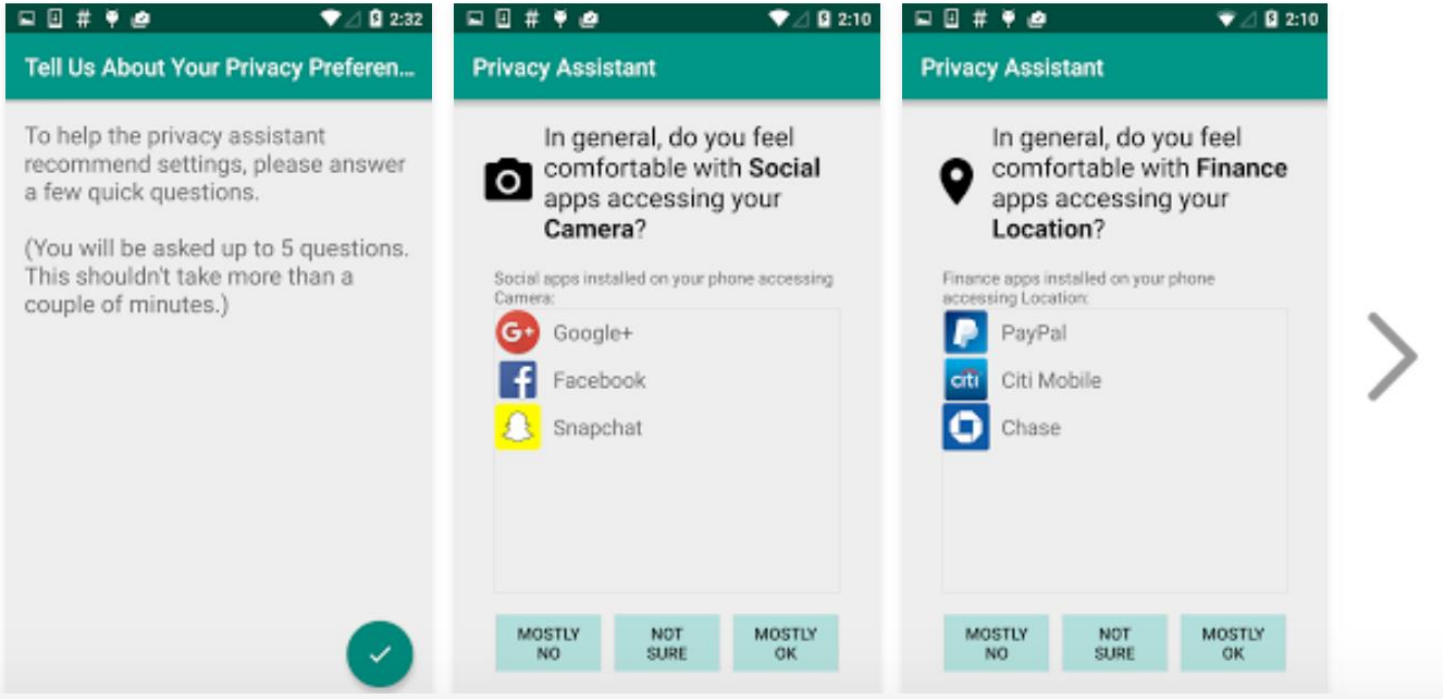

IoT: Additional Challenges

- No App Stores
- No (standardized) UI
- Often hidden, embedded
- More ways of collecting personal data
- Explosion in number of devices & services (scale)

Vision: Personalized Privacy Assistants (PPAs)

- **Selectively notify** us about privacy practices we may **not be expecting, yet care about**
- **Learn** many of our privacy preferences and expectations, and **semi-automatically configure** many settings on our behalf
  - Able to explain their recommendations/decisions
- **Learn** the best way of communicating with us based on the context at hand
- **Nudge us** to occasionally revisit/more carefully consider some of our preferences and decisions
- The assistants should ideally work across all domains and be **minimally disruptive**

Overall Architecture*

Policy Enforcement Points (PEP)
- Stores resource-specific and user-specific privacy policy settings
- Enforces settings for data collected by IoT resources

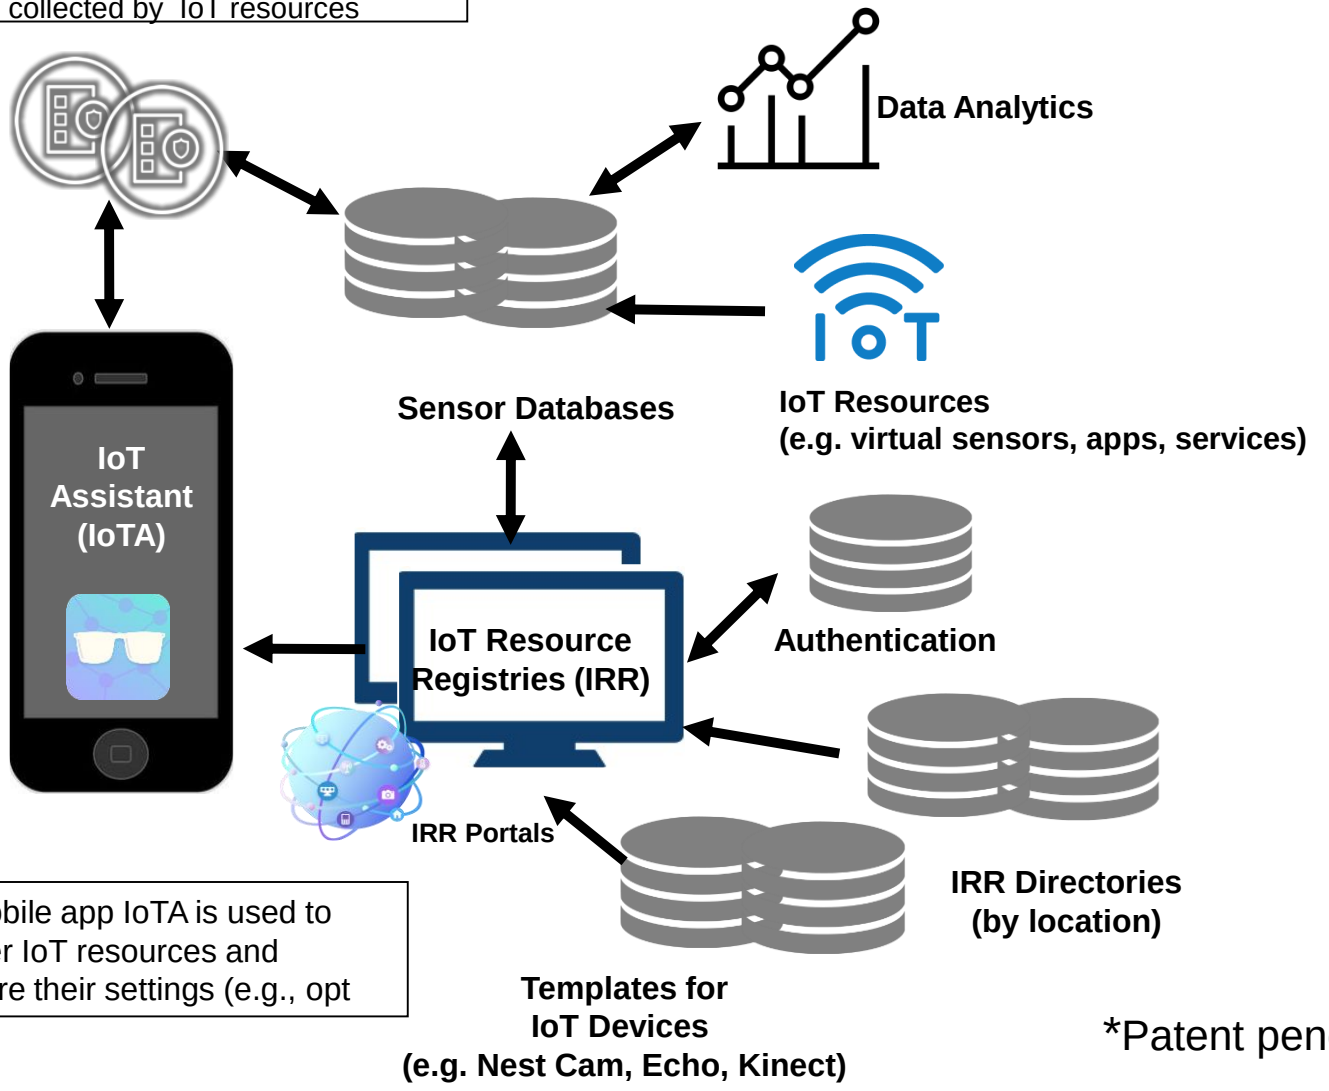

The mobile app IoTA is used to discover IoT resources and configure their settings (e.g., opt out)

*Patent pending

Deployment Example

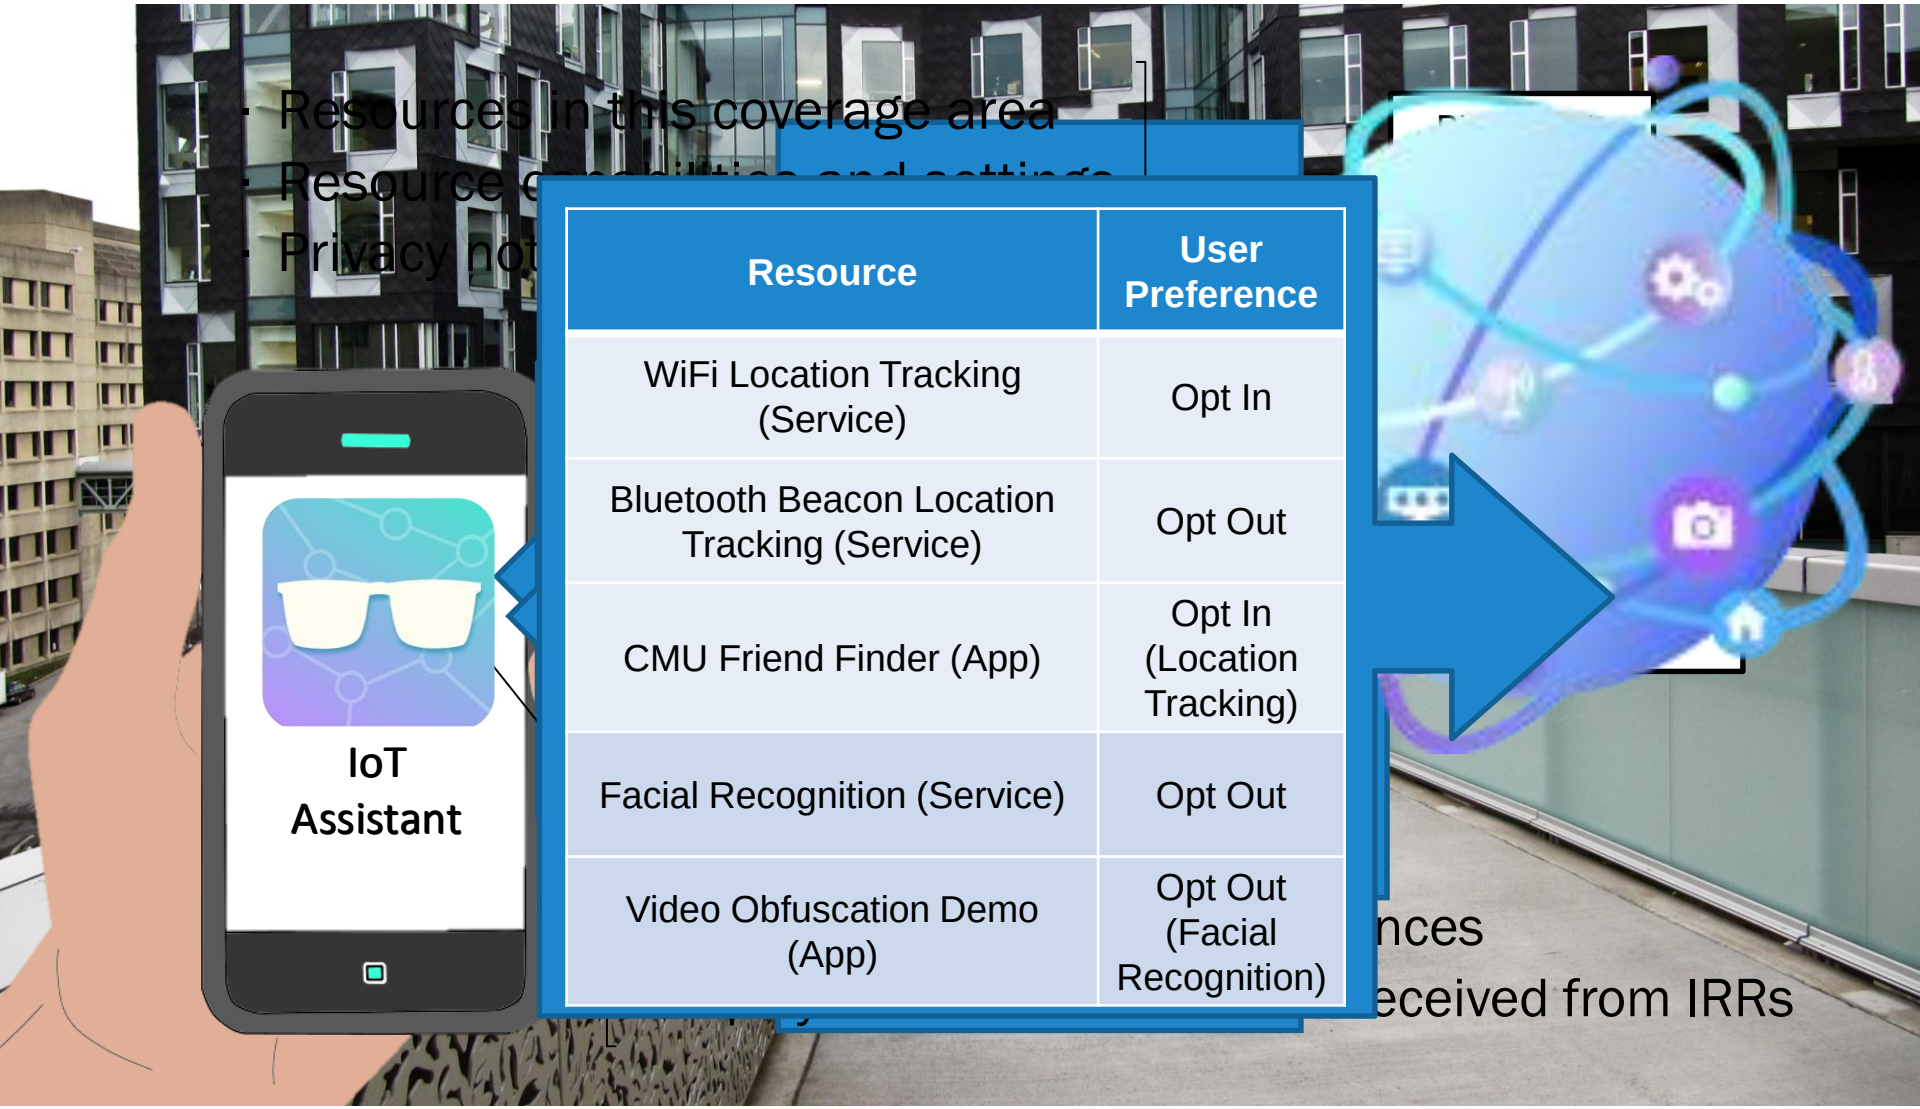

Resource	User Preference
WiFi Location Tracking (Service)	Opt In
Bluetooth Beacon Location Tracking (Service)	Opt Out
CMU Friend Finder (App)	Opt In (Location Tracking)
Facial Recognition (Service)	Opt Out
Video Obfuscation Demo (App)	Opt Out (Facial Recognition)

IoT Resource Registry Portal

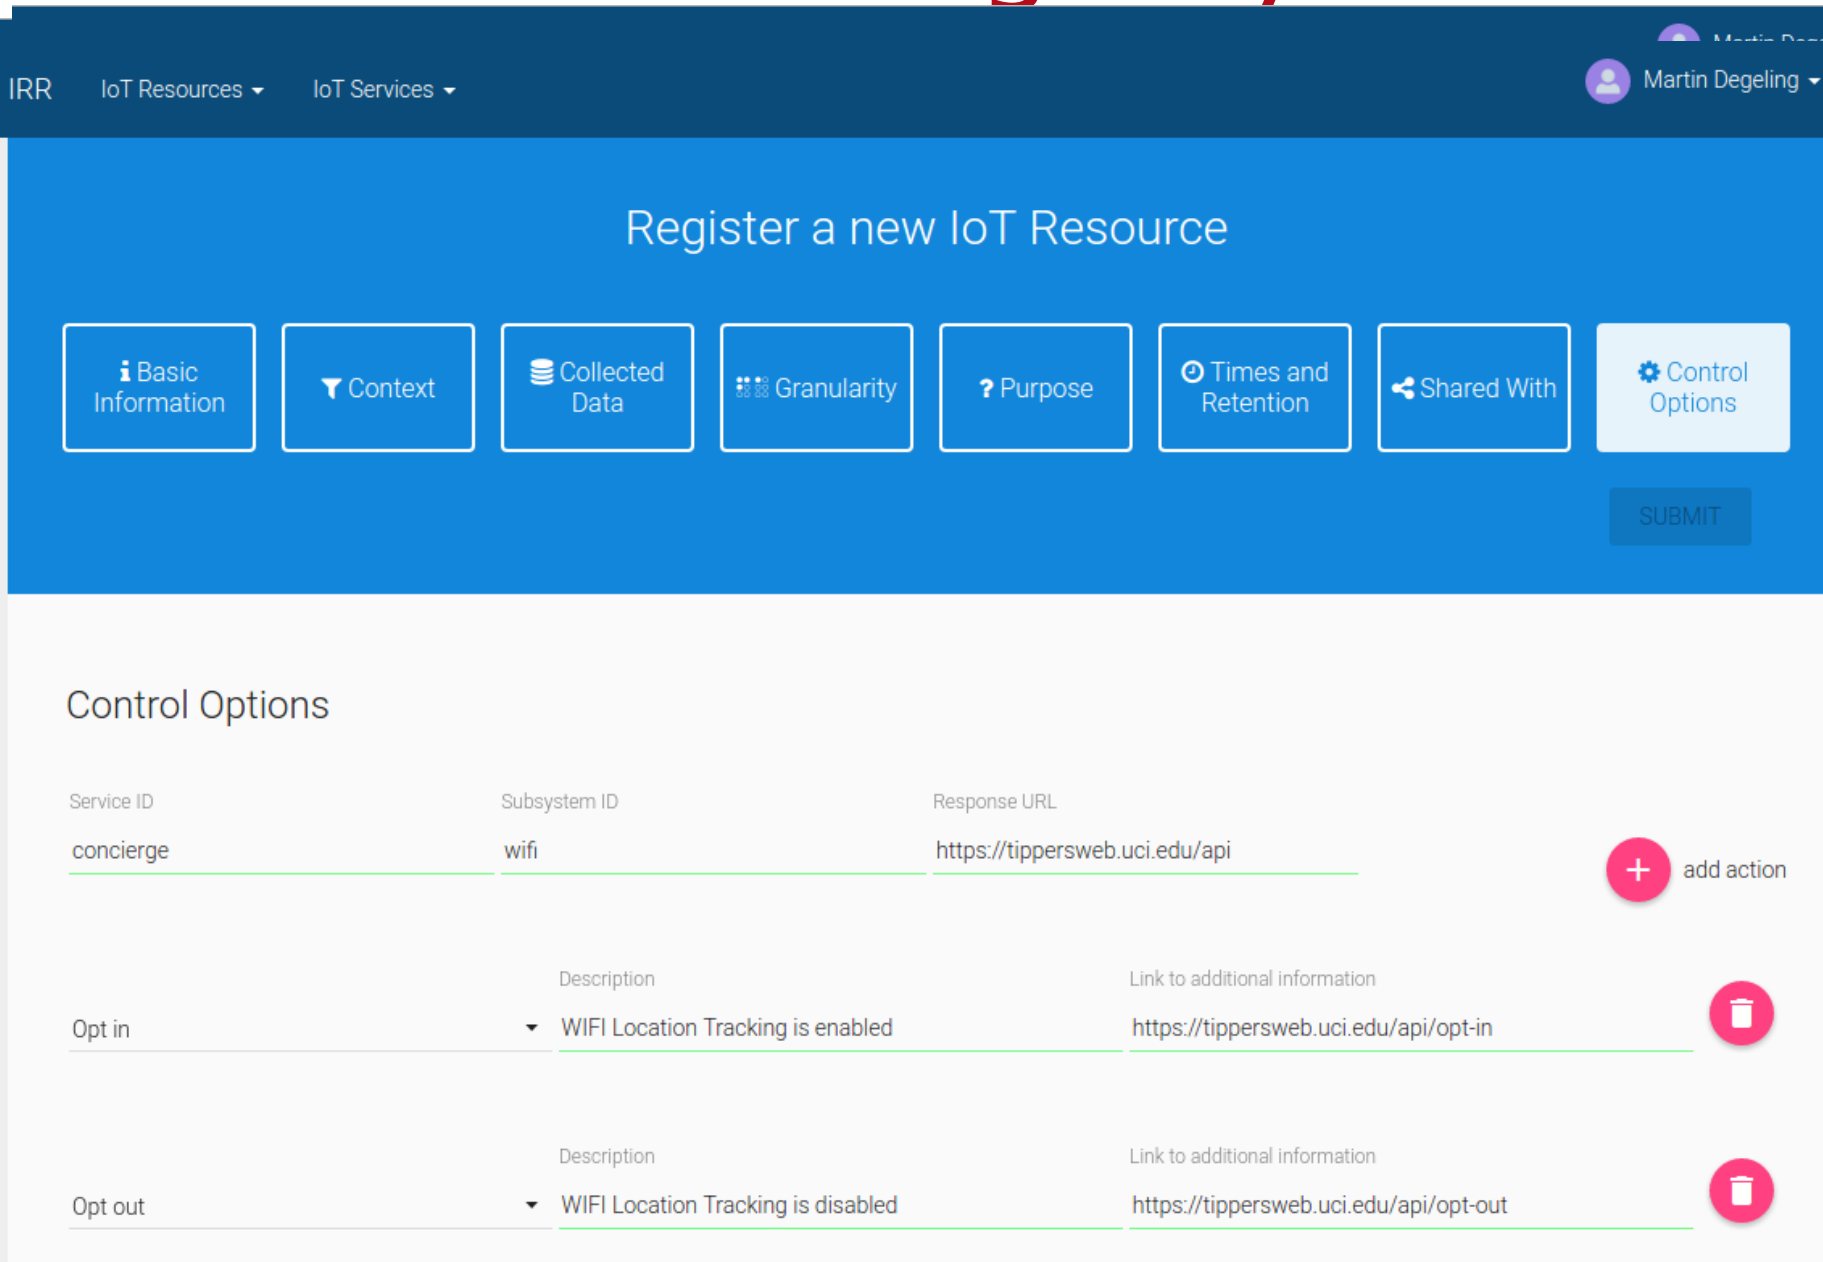

IoT Assistant

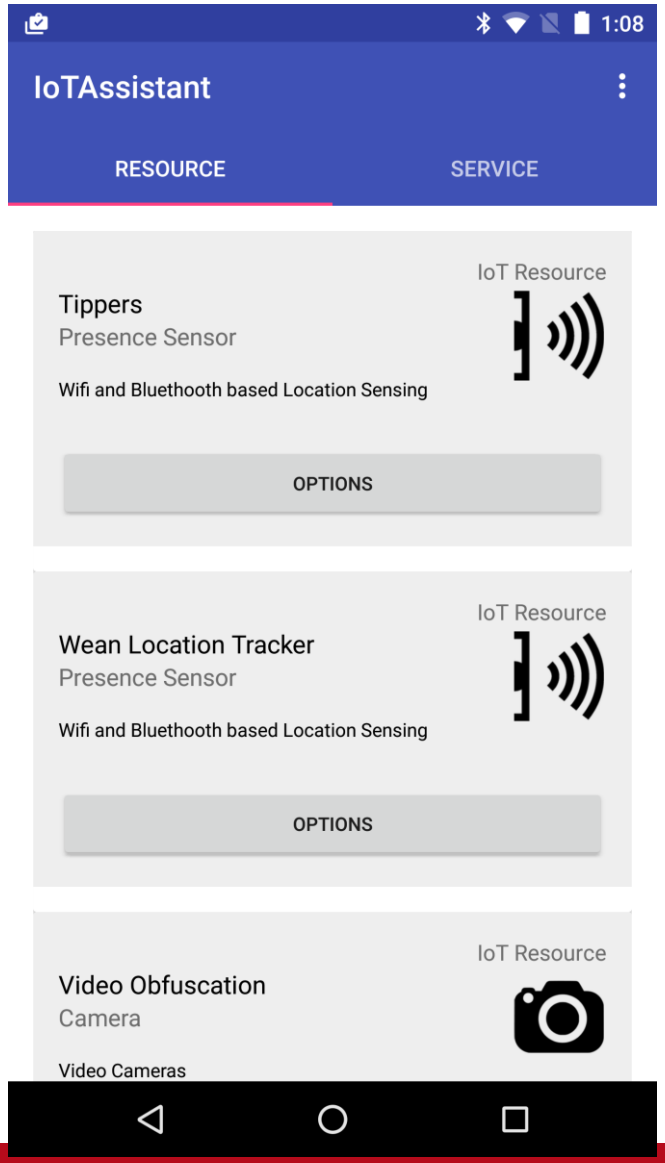

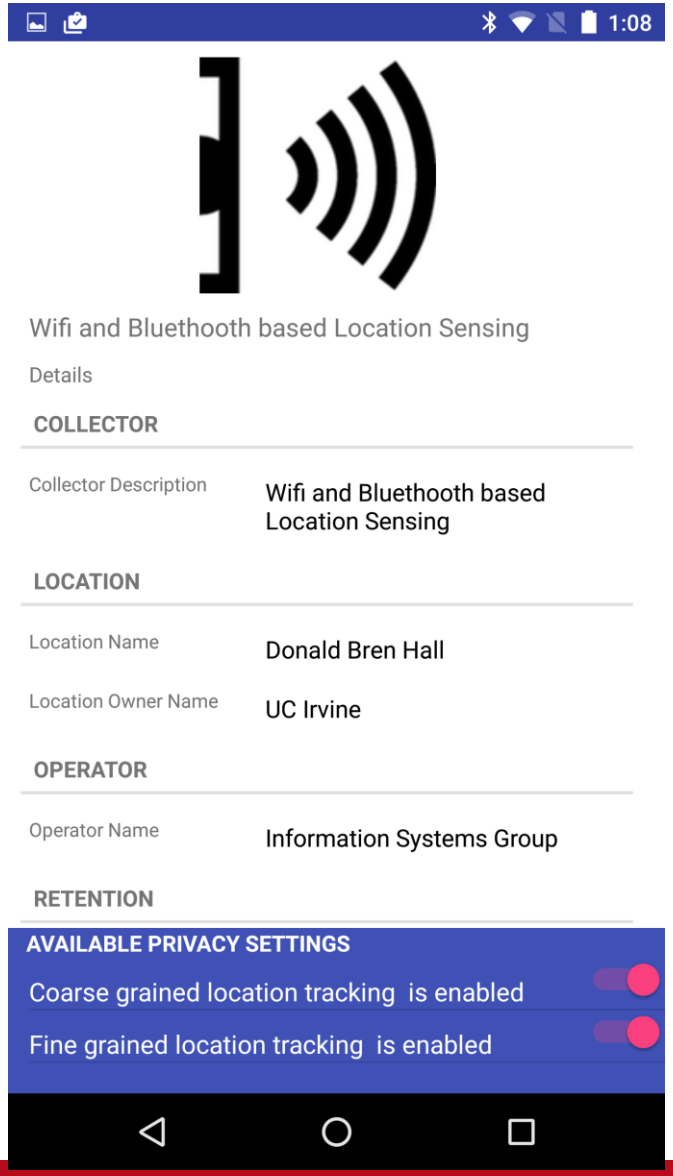

Possible Applications

- Smart buildings
- Smart cities
- Smart homes
- Smart cars
- etc.

Current Status

- Deployed at UC Irvine
- Deployed at CMU
- First public release coming out this summer

Privacy-aware Video Streaming

Train Facial Features

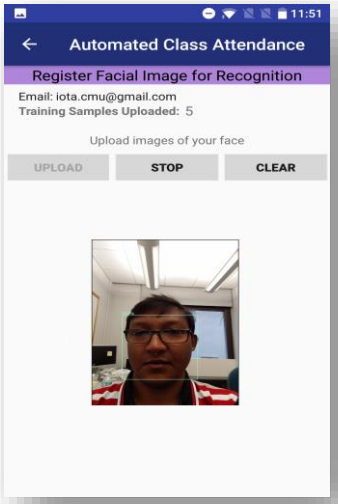

Control Opt-in

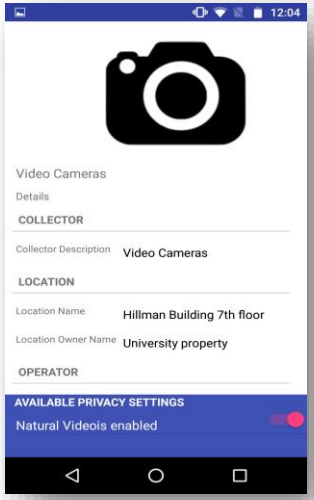

Live Video Stream

Monitor Class Attendance

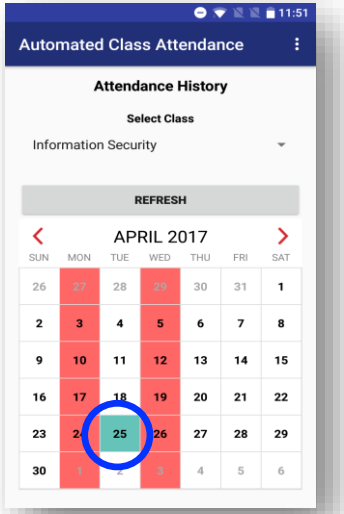

Demo: https://goo.gl/gtpbpK

IoT Device Templates

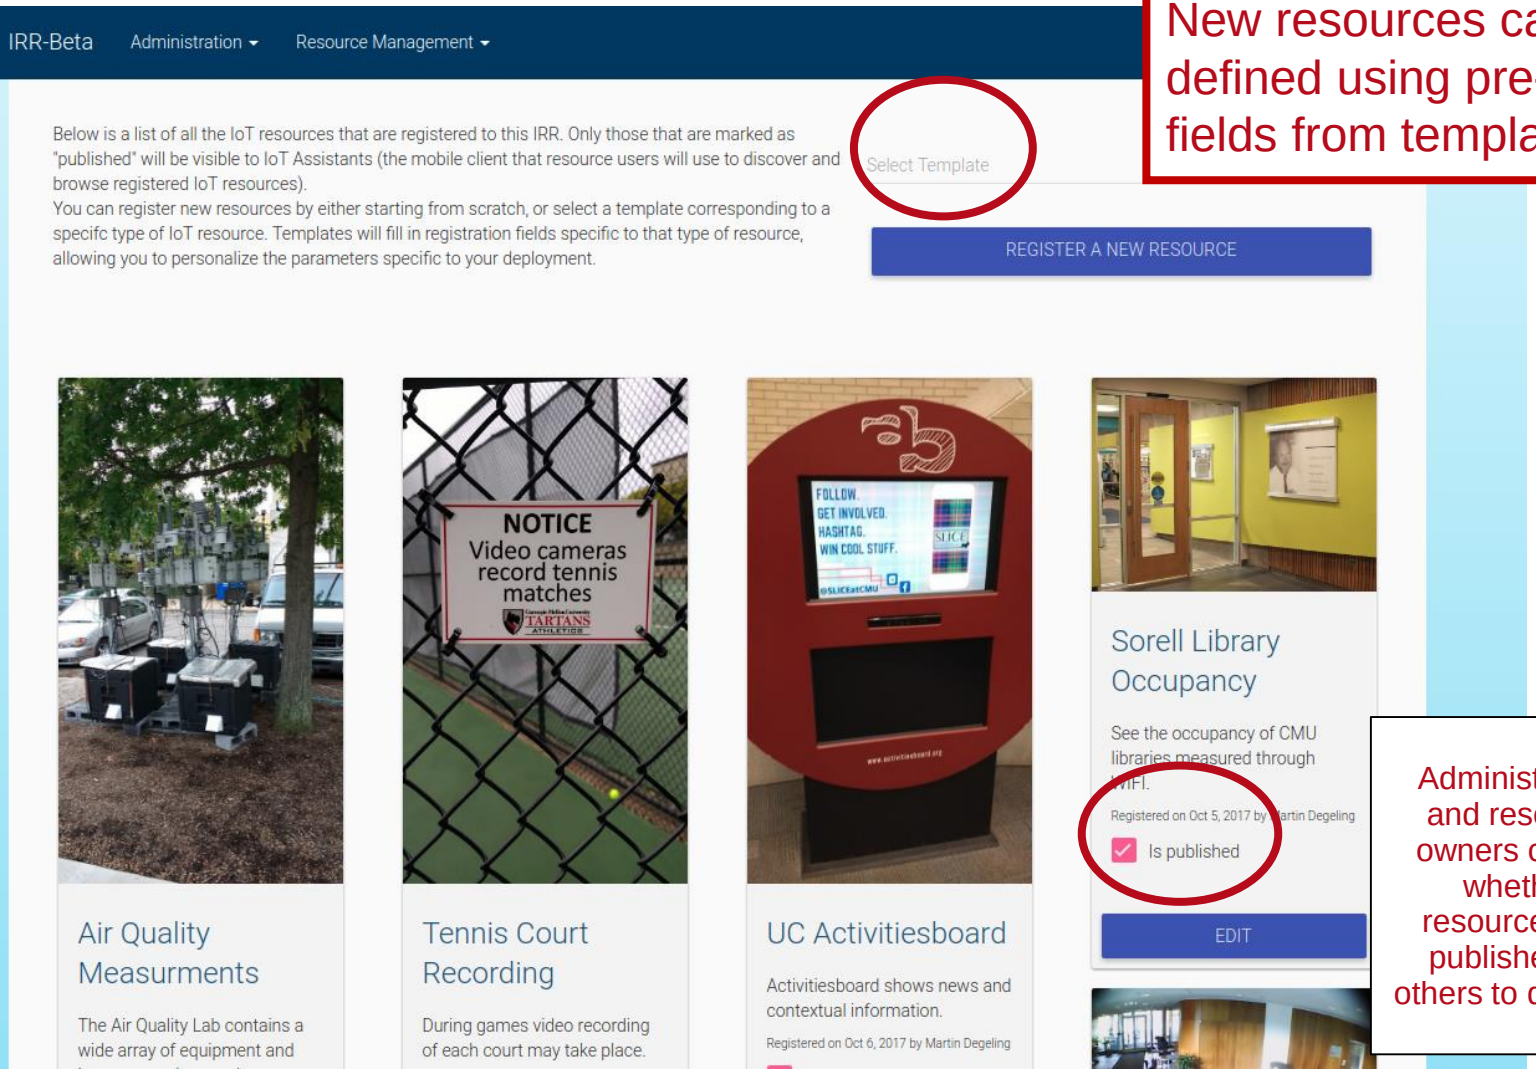

New resources can be defined using pre-filled fields from templates

Existing Templates:
- Echo
- Honeywell thermostat
- Kinect
- Google Home
- Cujo
- Nest Cam
- Wink Relay
- Zensors

Administrators and resource owners control whether resources are published for others to discover

Concluding Remarks - I

- AI and IoT are already widely deployed today
- They open the door to increasingly rich data collection and use processes and offer many promises
- But, if left unchecked, they are also a clear threat to our expectations of privacy

Concluding Remarks - II

- Regulations such as the Hong Kong Personal Data Ordinance or the new EU GDPR are establishing standards intended to protect our privacy
  - GDPR is the first regulation to come with steep penalties and it has global reach
  - But **regulations by themselves are not sufficient**
- These regulations are calling for the development of new technologies.
- **AI can be expected to play an important role** in addressing the **scale and complexity** of privacy issues that have emerged… **as a result of ….AI and IoT**

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

The **Usable Privacy Policy Project** and the **Personalized Privacy Assistant Project both** involve a collaborations with a number of individuals.
See **usableprivacy.org** and **privacyassistant.org** for additional details incl. lists of collaborators and publications

Q&A

Backup Slides

Hong Kong Personal Data Ordinance (Dec. 1996)

Six Principles:

1. **Purpose & Manner of Collection** has to be disclosed to data subject
2. **Accuracy and Duration of Retention** of Personal Data: data should be uptodate and only retained as long as necessary
3. **Use of Personal Data**: only for the purpose for which data was collected – unless otherwise agreed by data subject
4. **Security of Personal Data**: protection against unauthorized or accidental access, processing or deletion
5. **Notification**: Open policies about data being collected & for what purpose
6. **Access to personal data**: right to review and correct data about oneself

Hong Kong Personal Data Ordinance

- Personal data can only be used for the **purpose** for which it was collected – **no frivolous collection**
  - This also restricts sharing
- Purpose has to be **stated from the beginning**
- People should have the right to inspect information held about them within 40 days of their asking
  - May involve a fee
- Data has to be **corrected if erroneous**
- Data has to be **secure**
- No direct marketing or teleselling **if someone opts out**
- Individuals can sue if damage results from the release of confidential data, or from inaccurate data or other breach
- **Note**: This is a very **approximate summary** – read the text of the Ordinance for a more detailed & accurate understanding

Recent Amendments to HK Personal Data Ordinance

- Effective April 1, 2013
- Imposes additional requirements for data users that seek to:
  - Sell personal data
  - Use personal data for their own direct marketing purposes
  - Provide personal data to another person for that other person's direct marketing purposes

More Details

香港個人資料私隱專員公署
Office of the Privacy Commissioner
for Personal Data, Hong Kong

Exercising Your Right of Consent to and Opt-out from Direct Marketing Activities
under the Personal Data (Privacy) Ordinance[1]

It is common for members of the public to receive unsolicited [illegible] calls, mail, email, messages and so on from direct marketers promoting various products and services.

Under the Personal Data (Privacy) Ordinance ("the Ordinance") organisations are required to notify you and obtain your consent before using your personal data in their own direct marketing activities or transferring the data to another person for use in the latter's direct marketing activities.

Despite your consent to use your personal data in direct marketing, direct marketers must notify you of your opt-out right when using your personal data in this manner *for the first time*. On the other hand, you may require them to cease to so use the data *at any time*. The request must be complied with at no cost to you. Further, despite your consent for an organisation to transfer your personal data to third parties for use in the latter's direct marketing activities, you may *at any time* require the organisation to cease to transfer the data and to notify any person to whom your personal data has been so transferred to cease to use the data in direct marketing. Again the request must be complied with at no cost to you.

For contraventions of the requirements under the Ordinance involving the transfer of personal data to third parties for gain, the maximum penalty is a fine of HK$1,000,000 and imprisonment for 5 years. For other direct marketing contraventions, the maximum penalty is a fine of HK$500,000 and imprisonment for 3 years.

The purpose of this leaflet is to explain the direct marketers' obligations when using your personal data and how you can exercise your right to indicate your consent to the intended use or transfer of your personal data in direct marketing. It also guides you to make an opt-out request under the Ordinance in order to effectively stop an organisation from continuing to use or transfer your personal data for direct marketing purposes.

EU – GDPR

- Took effect May 25, 2018
- Stricter provisions than EU Data Protection Directive
  - Single set of rules and **one-stop shop model**: each company coordinates with a single Supervisory Authority
  - **Privacy by Design and by Default** – incl. default privacy settings that are protective
  - Opt-in: data controllers must be able to **prove consent & consent may be withdrawn**
  - **Severe penalties** for violations: up to 20MEUR or 4% of worldwide turnover, whichever is greater
  - Right to **"erasure"**
  - Data **portability**