Illustration: Matt Stevens

Rethinking Access Control In the Home IoT

CMSC 23210/33210 Usable Security and Privacy

Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth,
Earlence Fernandes, Blase Ur

Back to the Future Part II - Universal Pictures - cnn.com


EXIT
c|net
CNET Smart Apartment - cnet.com

From Single User to Multi User

Single User

iPhone X, macOS High Sierra - apple.com

From Single User to Multi User

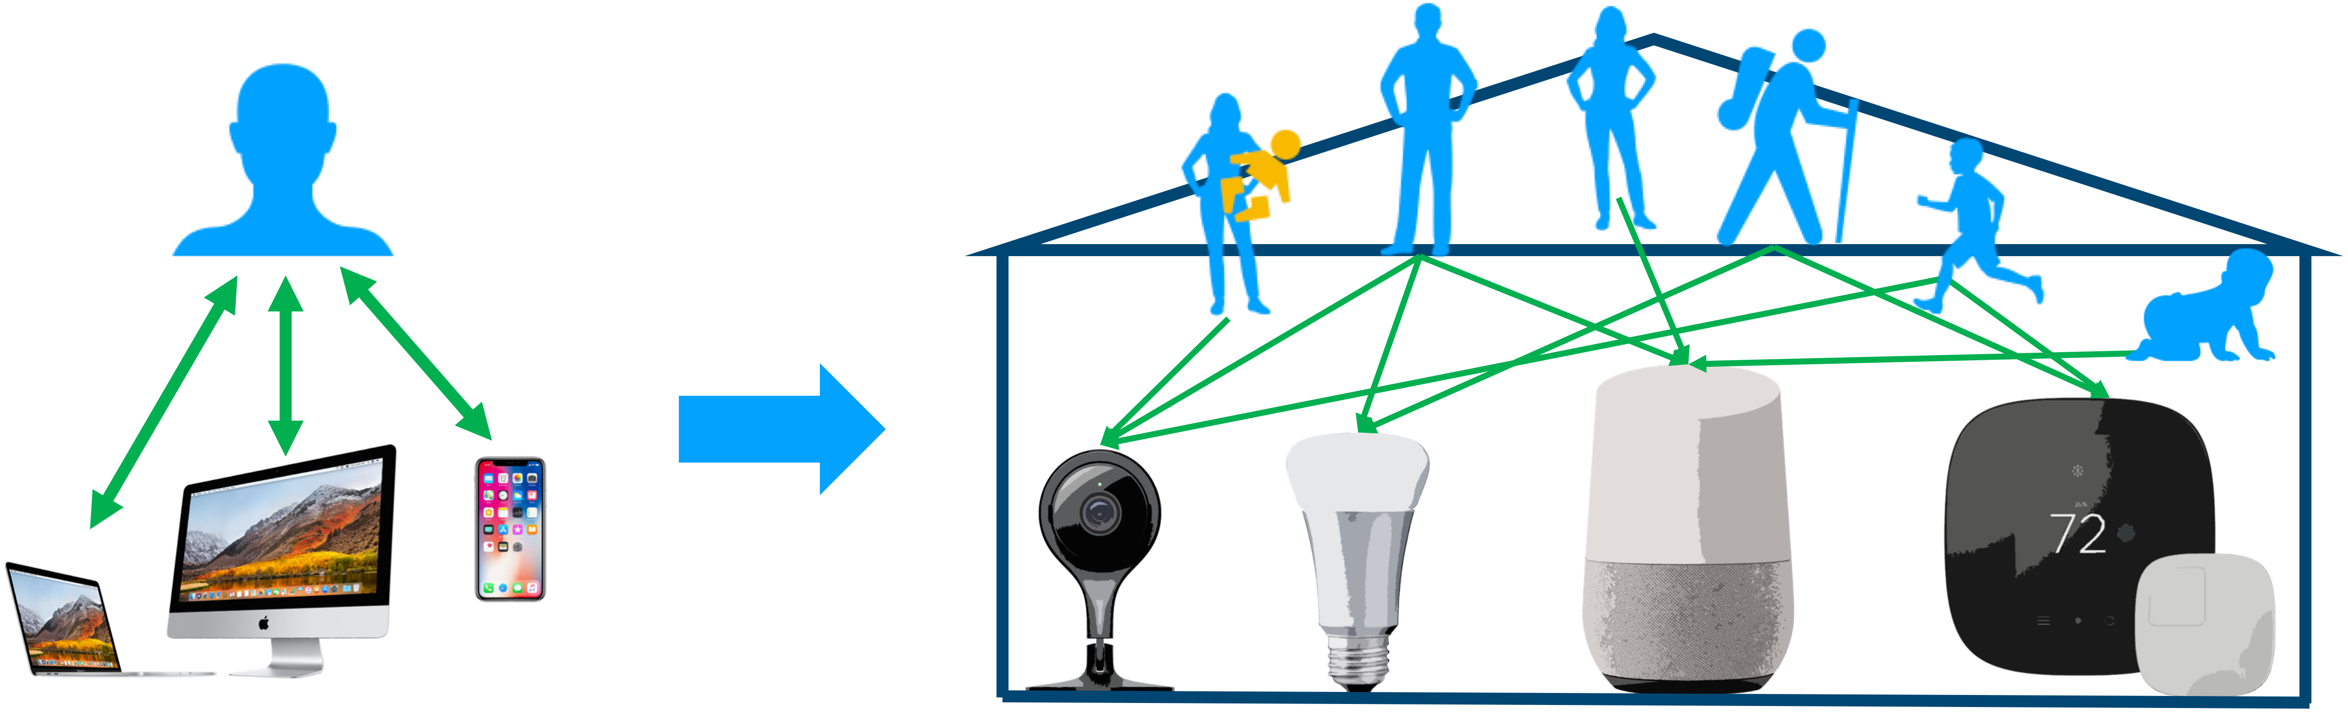

Single User

Multi User

Vendors Still Treat It The Old Way!

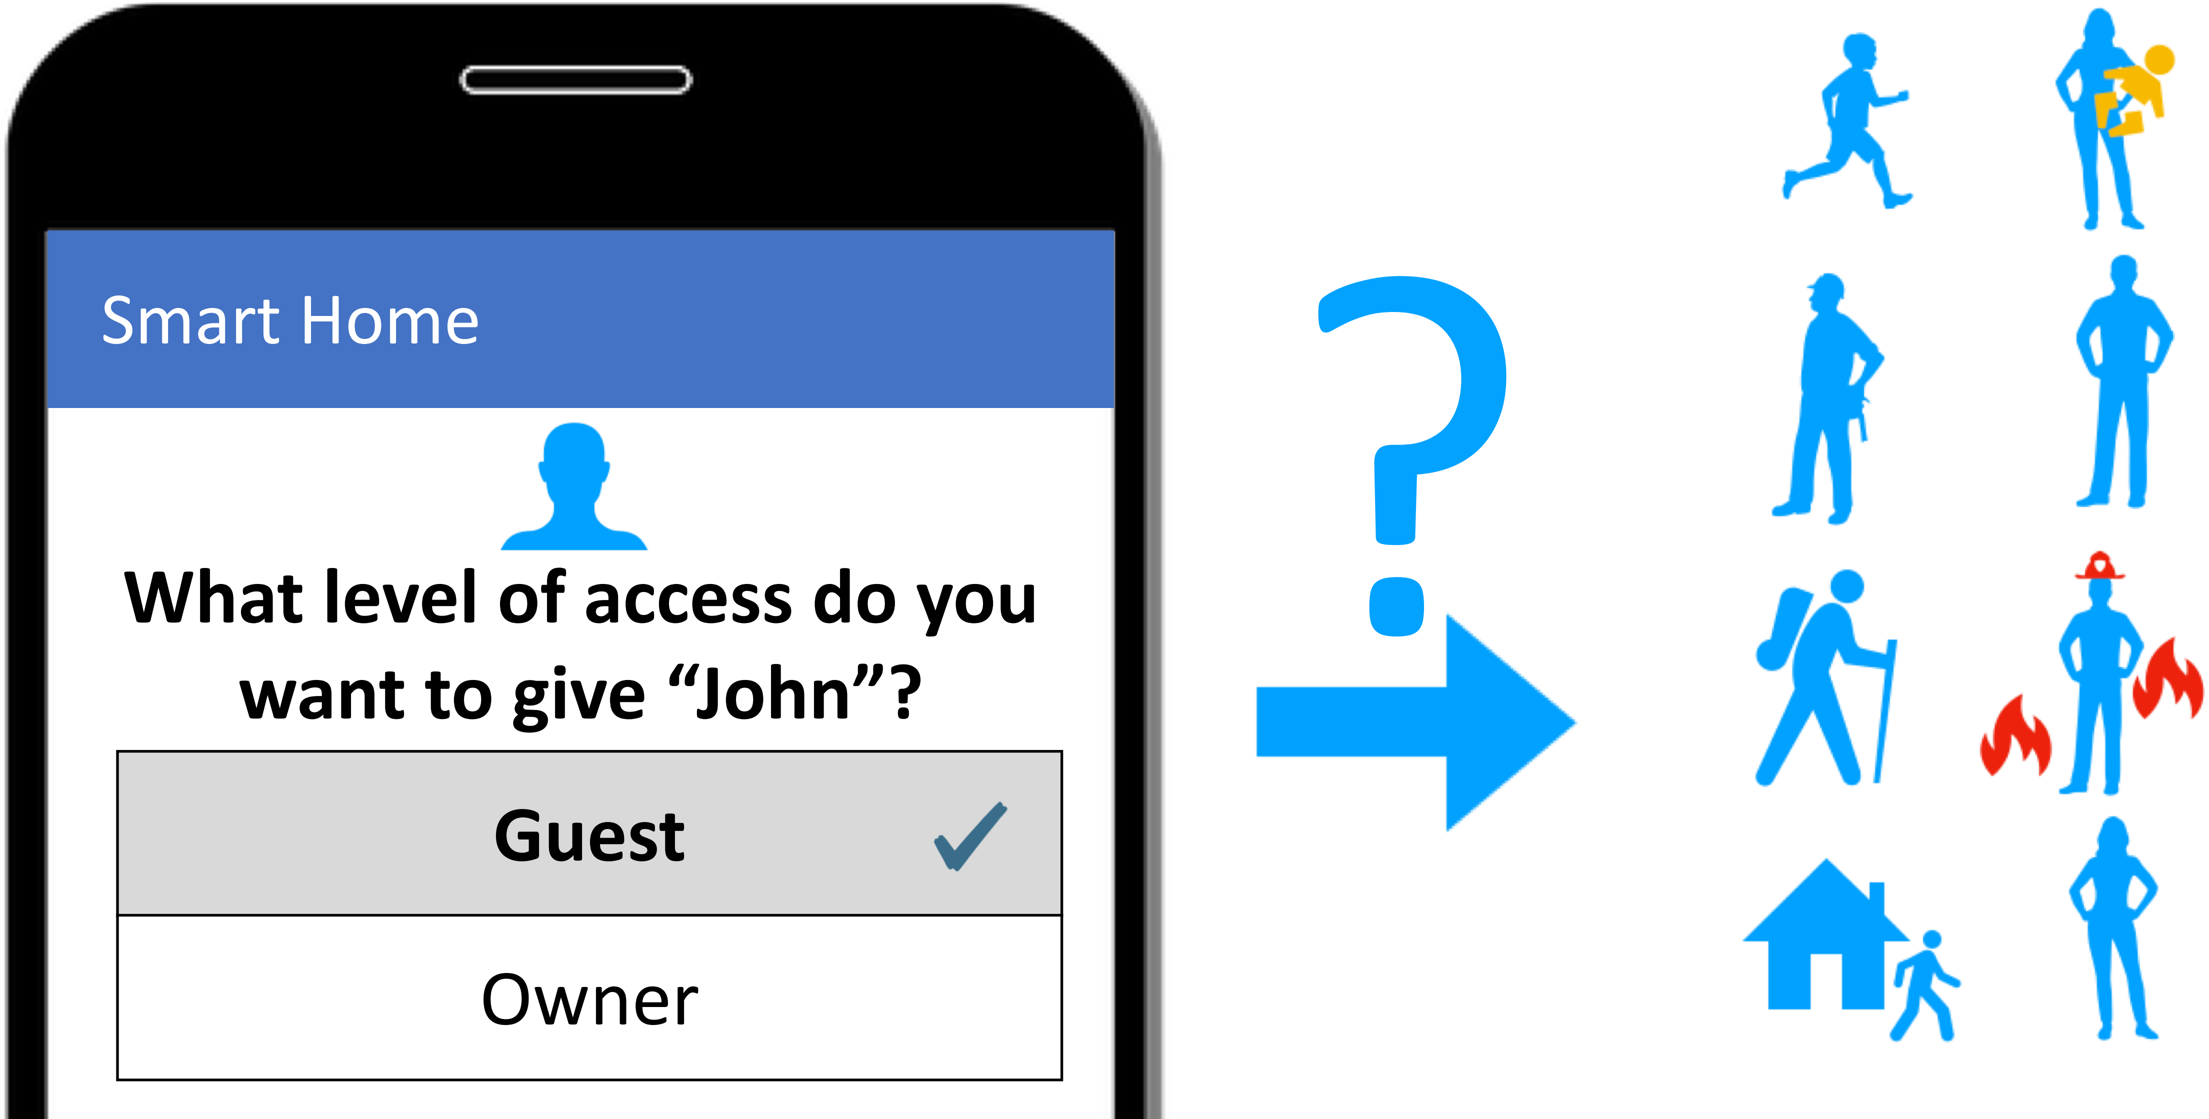

Please Enter Your Password:

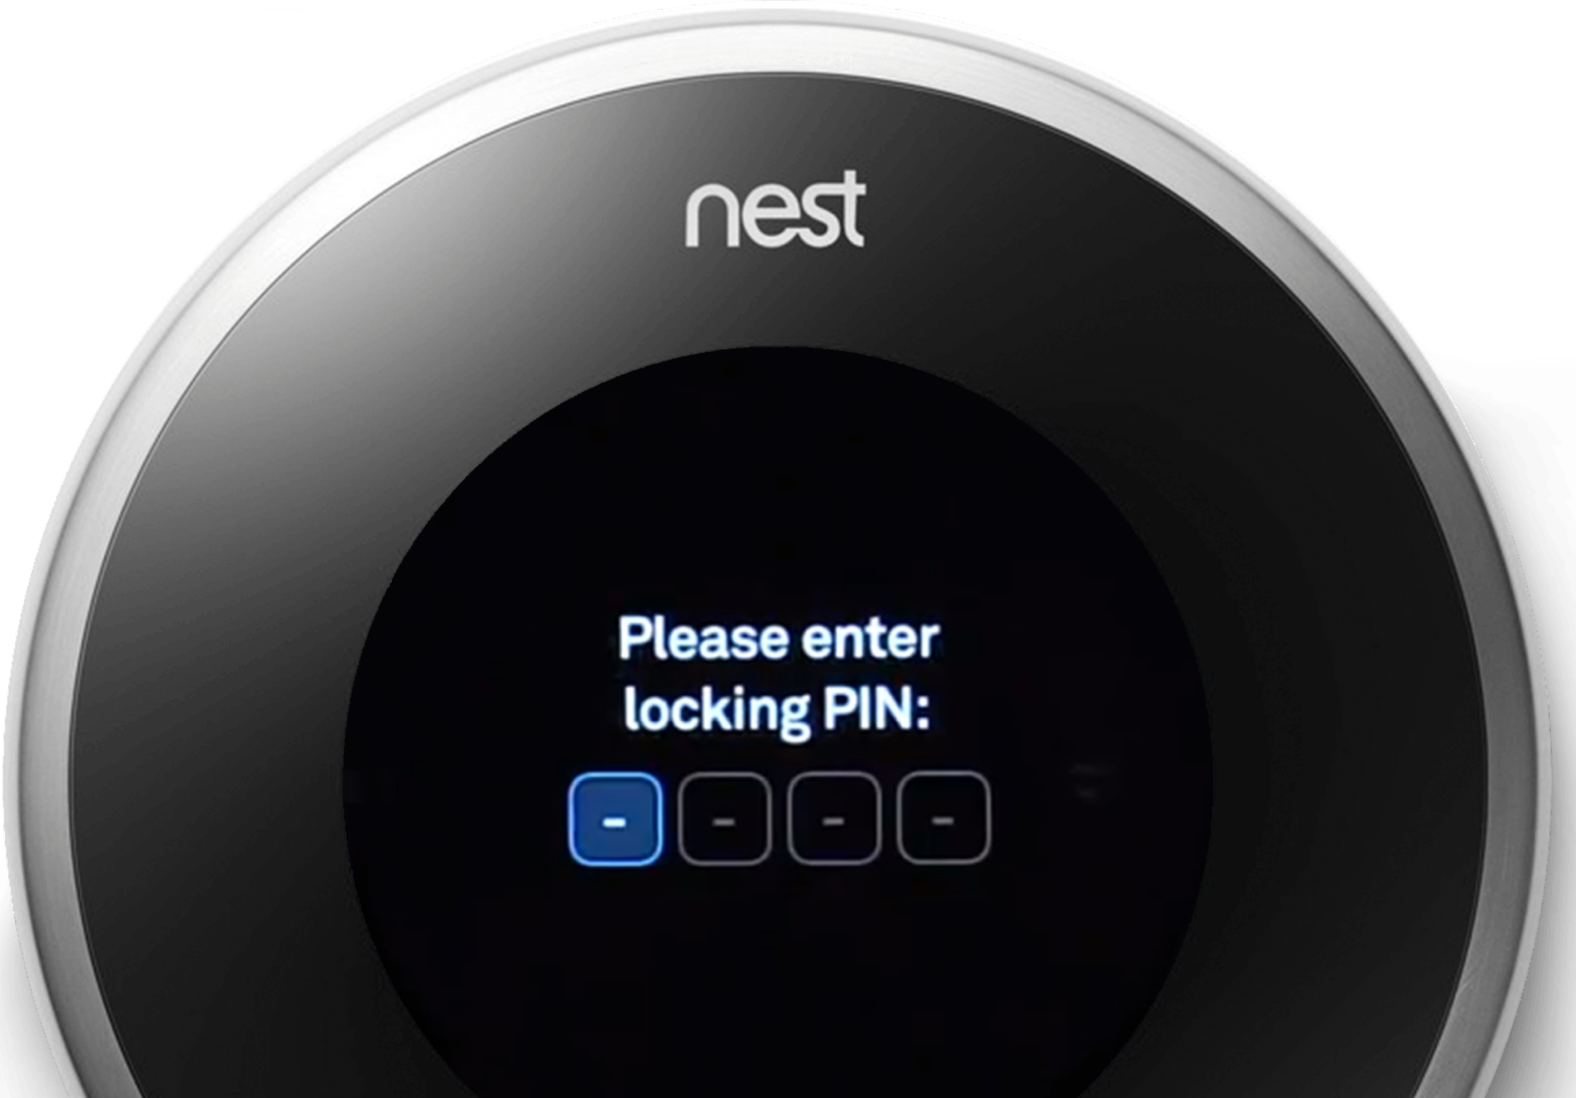

Nest Learning Thermostat
nest.com

Please Enter Your Password:

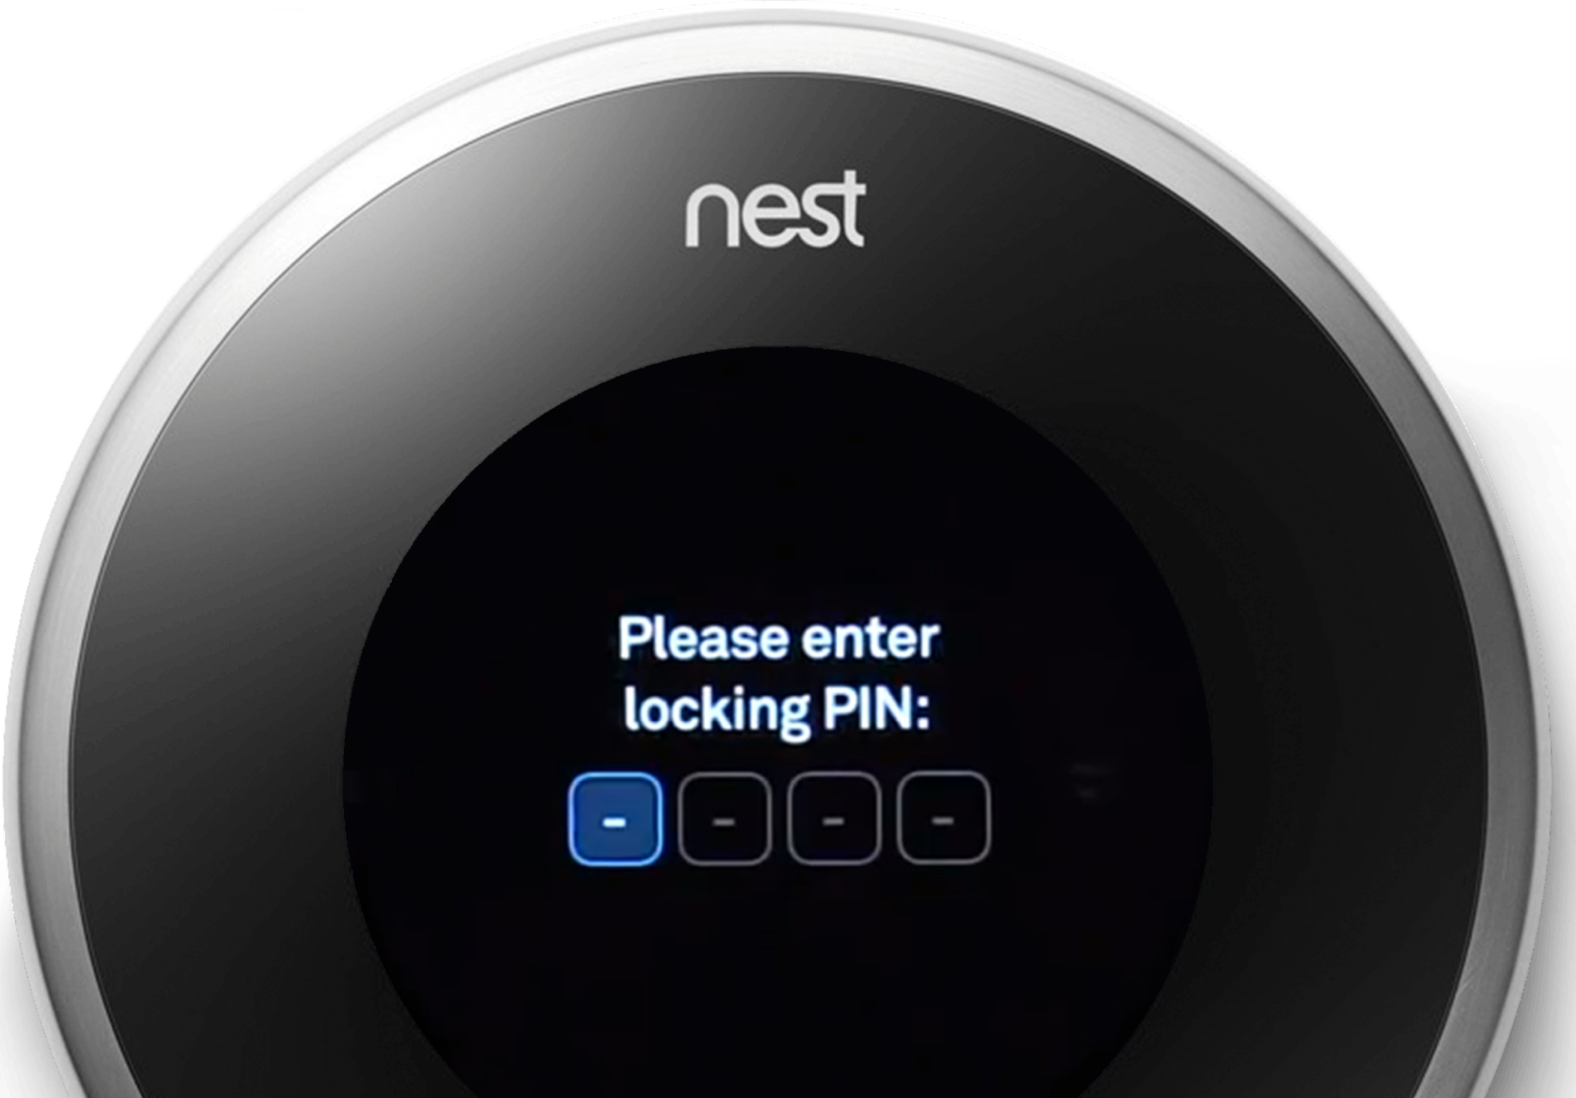

South Park Logo
deviantart.com

Home IoT Devices

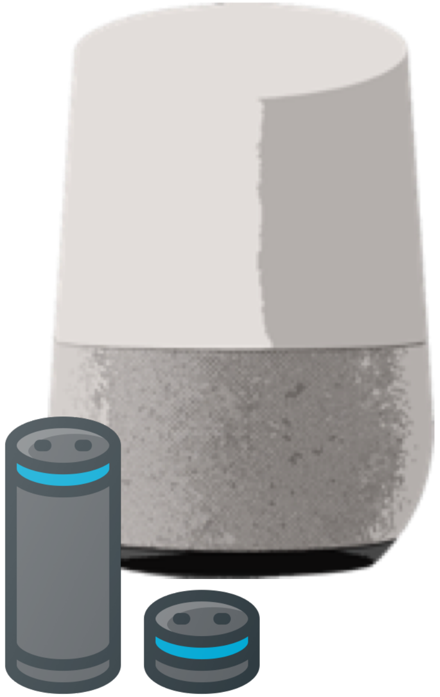

“Play music!”
“Order me a puppy!”

Icon made by Those Icons, Smashicons, Vectors Market from www.flaticon.com

Just to Summarize…

Traditional Devices	Home IoT
A Single User	Multiple Users
With Keyboard & Screen	Often Without Keyboard & Screen
Device-level Access Control	Capability-level Access Control

Research Goals

We conducted a user study to…

- Map desired access-control policies for home IoT devices
    - How policies vary by relationships and capabilities
    - Identify potential default policies

Method

Before Implementing the Survey…

- What do relationships and capabilities mean for home IoT?

6 Relationships

24 Relationships

6 Relationships

- Your Spouse
- Your Teenage Child
- Your Child in Elementary School
- A Visiting Family Member
- The Babysitter
- Your Neighbor

Icons made by Vectors Market from www.flaticon.com

22 Capabilities

1)

2)

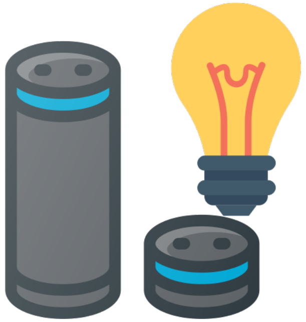

22 Capabilities

1)

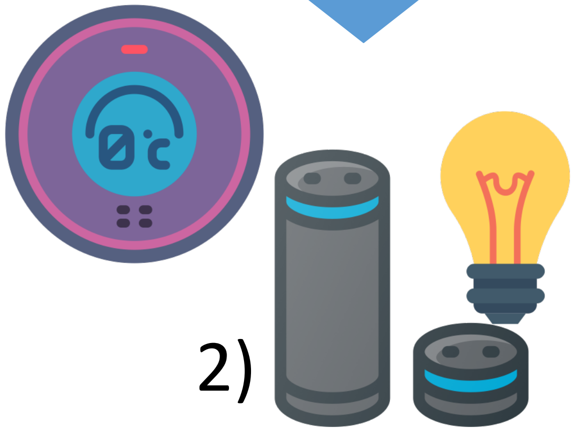

2)

3)

Order Online

Live Video

Answer Door

Mower Rule

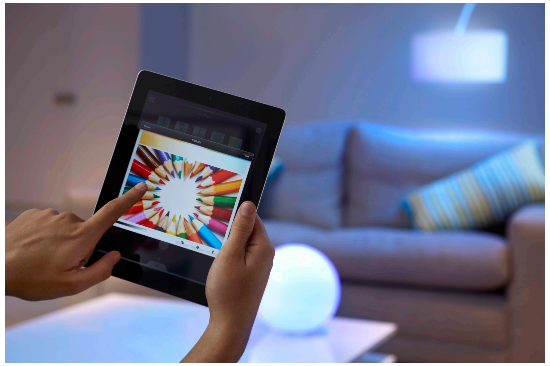

Lights Rule

Alexa - amazon.com, Nest Cam - nest.com, Ring Doorbell - thewirecutter.com, Husqvarna Automower - husqvarna.com, Hue - philips.com

User Study

Imagine you are the owner of a <smart device>.

Using this device, some users can access the following feature:

<capability>.

6x

When should <relationship> be able to use this feature?

- Always
- Sometimes
- Never

User Study

Imagine you are the owner of a Smart Voice Assistant.

Using this device, some users can access the following feature:

Make online purchases (e. g., on Amazon) on a shared household account.

When should your spouse be able to use this feature?

- Always
- Sometimes
- Never

Are Relationships and Capabilities Enough?

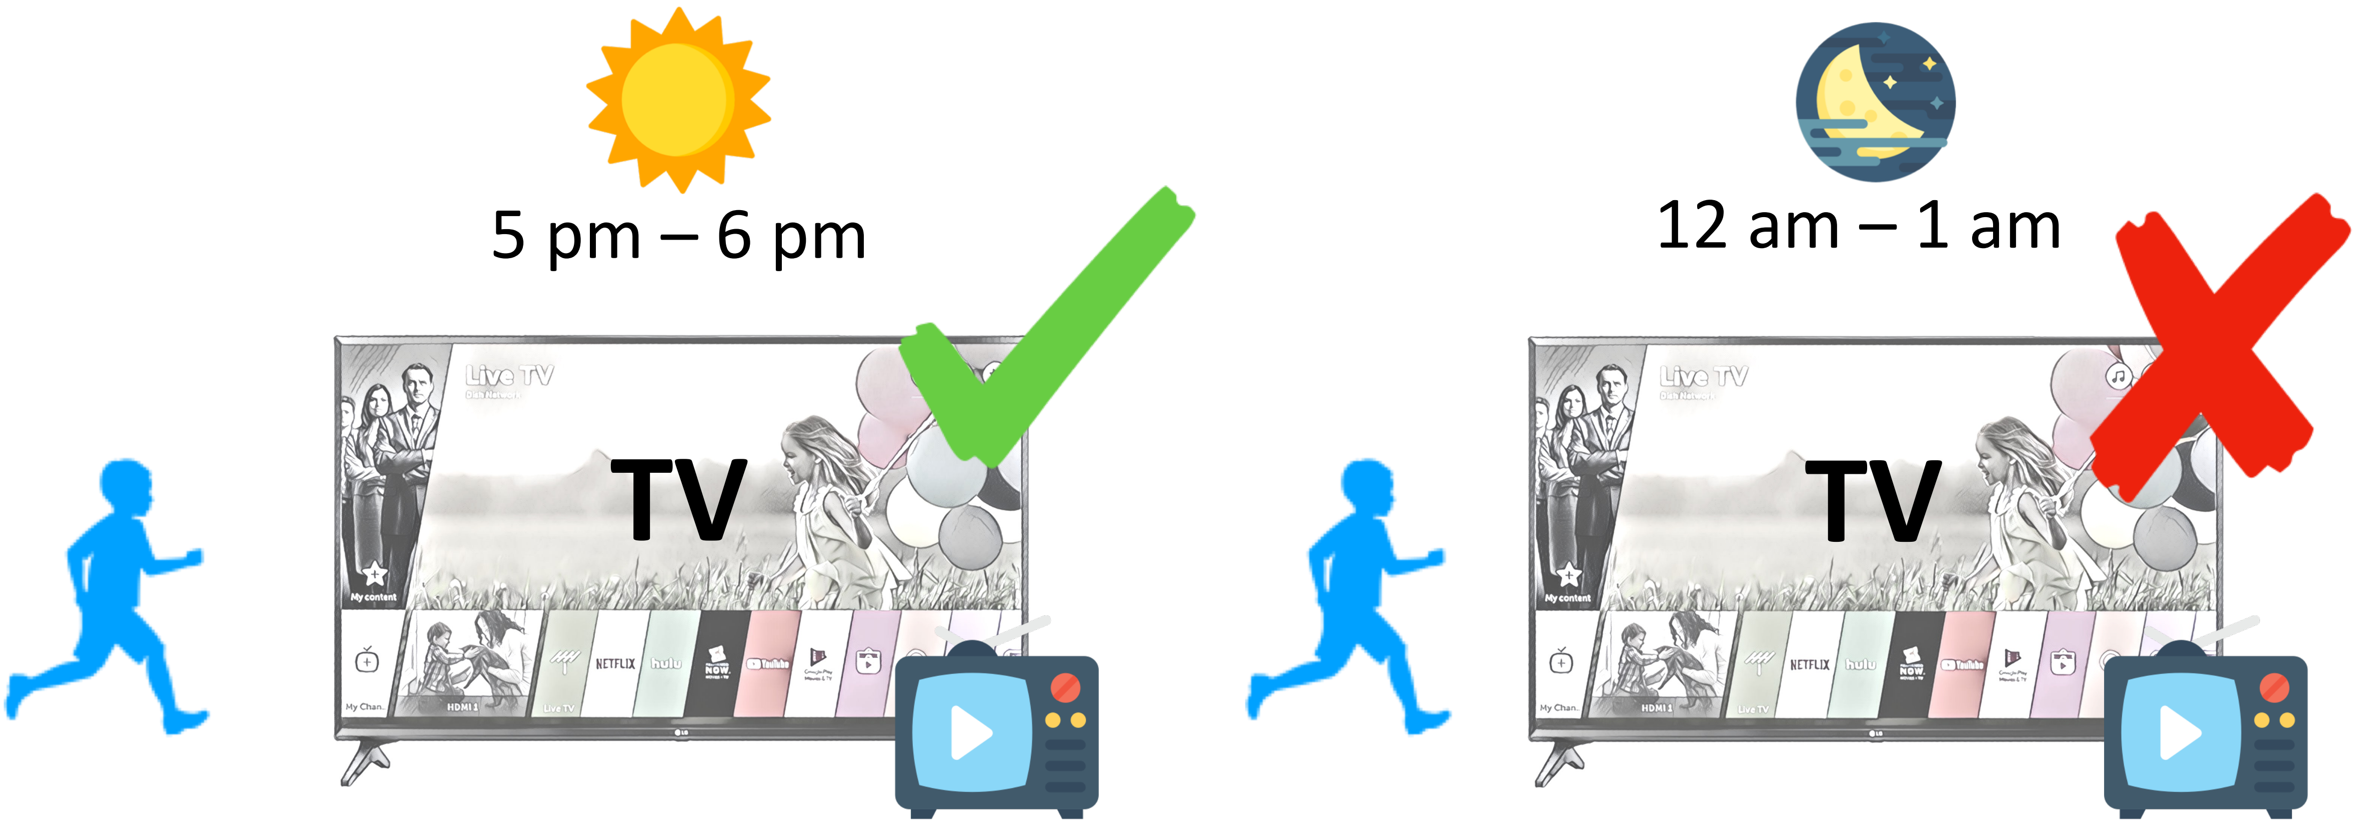

LG Smart TV – lg.com; Icons made by Freepik and Pixel Buddha - CC 3.0 BY

Research Goals

We conducted a user study to…

- Map desired access-control policies for Home IoT Devices
  - How policies vary by relationships and capabilities
  - Identify potential default policies
- What contextual factors affect the user's decision?

User Study

Imagine you are the owner of a Smart Voice Assistant.

Using this device, some users can access the following feature:

Make online purchases (e. g., on Amazon) on a shared household account.

When should your spouse be able to use this feature?

- Always
- Sometimes
- Never

User Study

- When should they have access to this capability?
- When should they **not** have access to this capability?

Results

425 Participants

54% Male

46% Female

Age 25-34 47%

CS 19%

Home IoT Device 44%

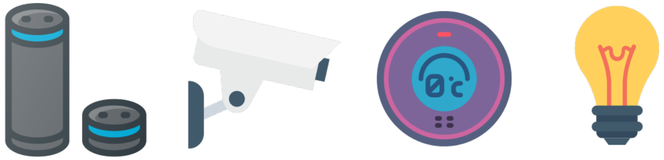

Results

Given one particular capability, what access-control policy should be set up for whom?

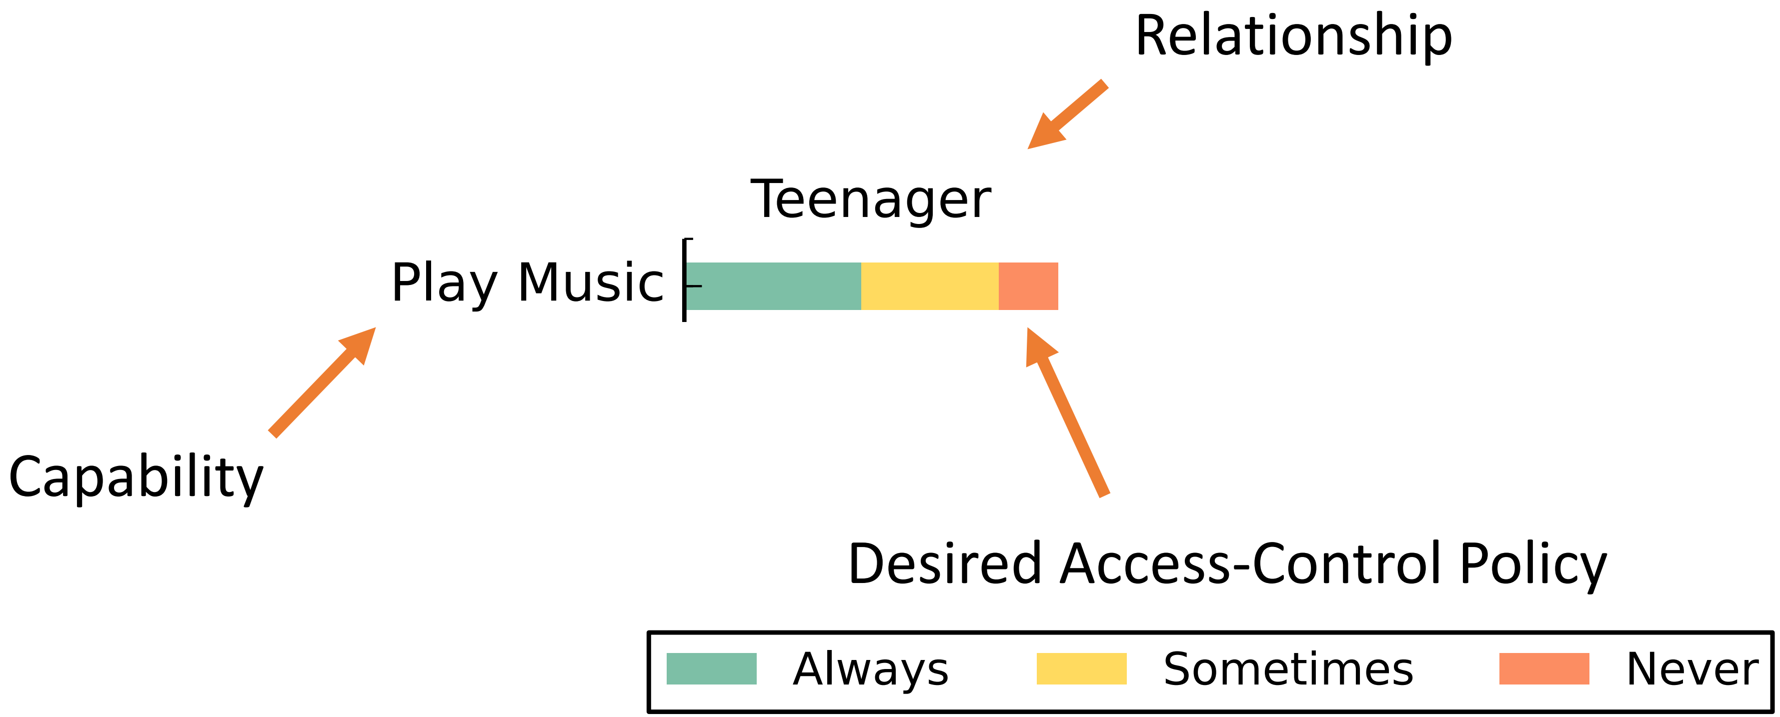

Comparison Between Capabilities

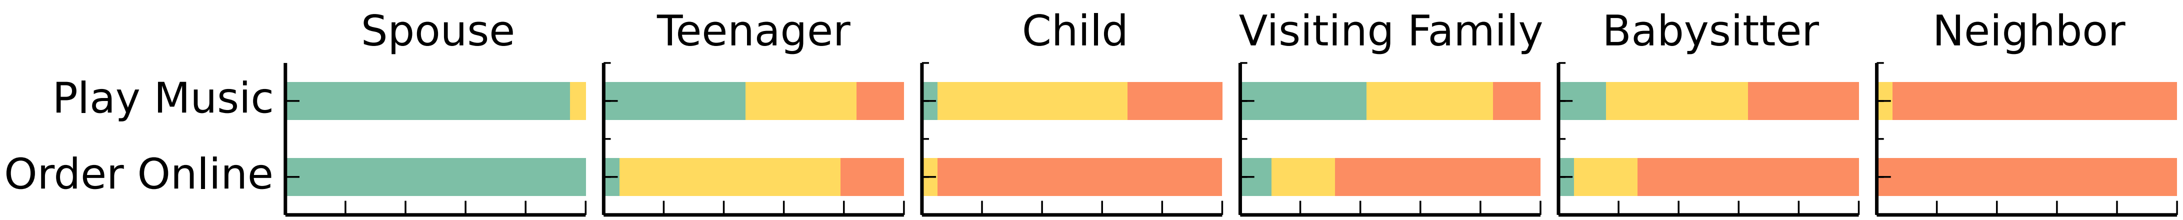

Icon made by Alexander Madyankin, Roman Shamin, Echo Plus with built-in Hub – Black - www.amazon.com

Capabilities Within One Device


Answer Doorbell


Delete Lock Log

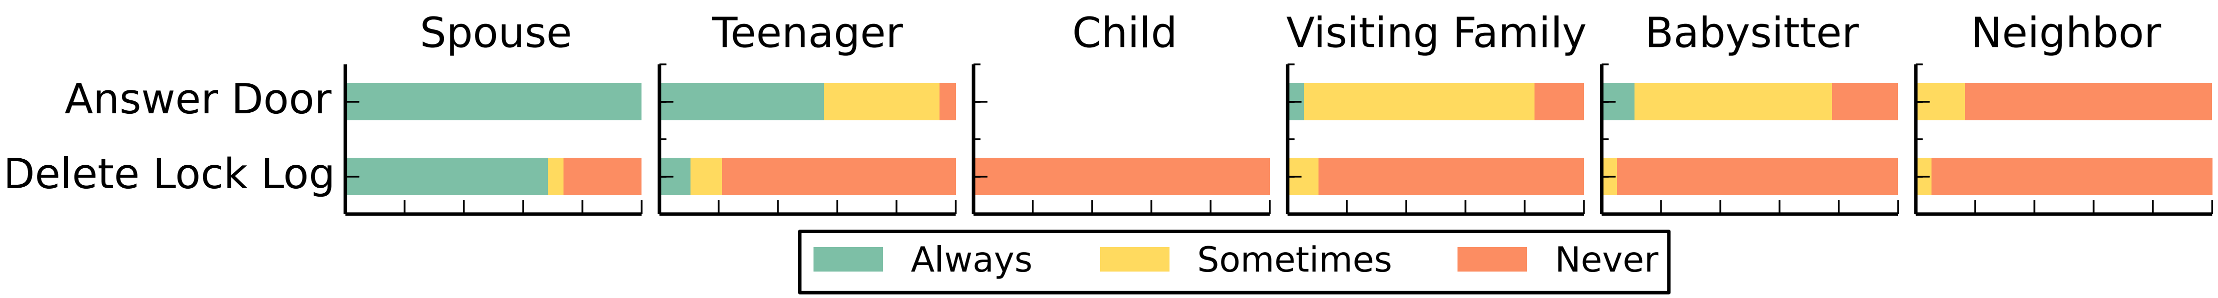

Amazon Key - www.amazon.com

Spouse Can Do Almost Everything

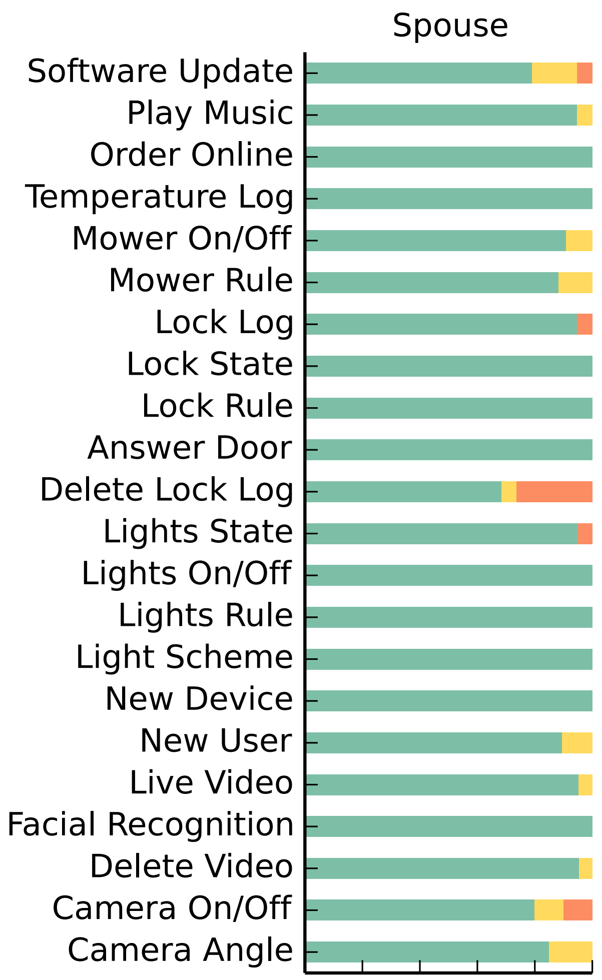

Neighbor Can Do Nothing

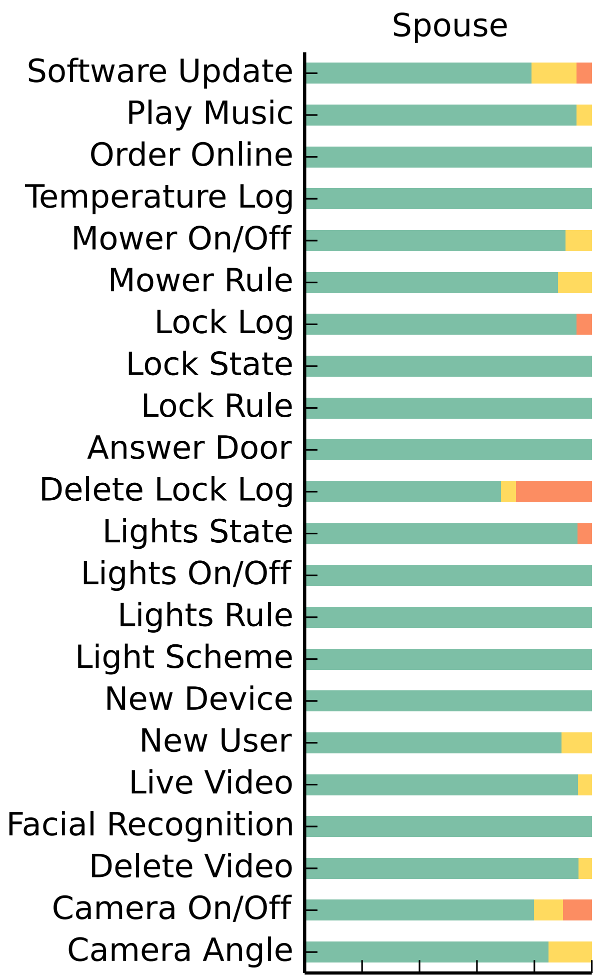

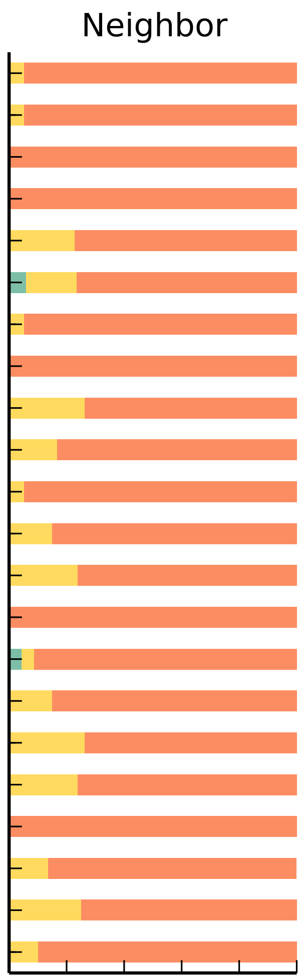

Other Relationships Are More Complex

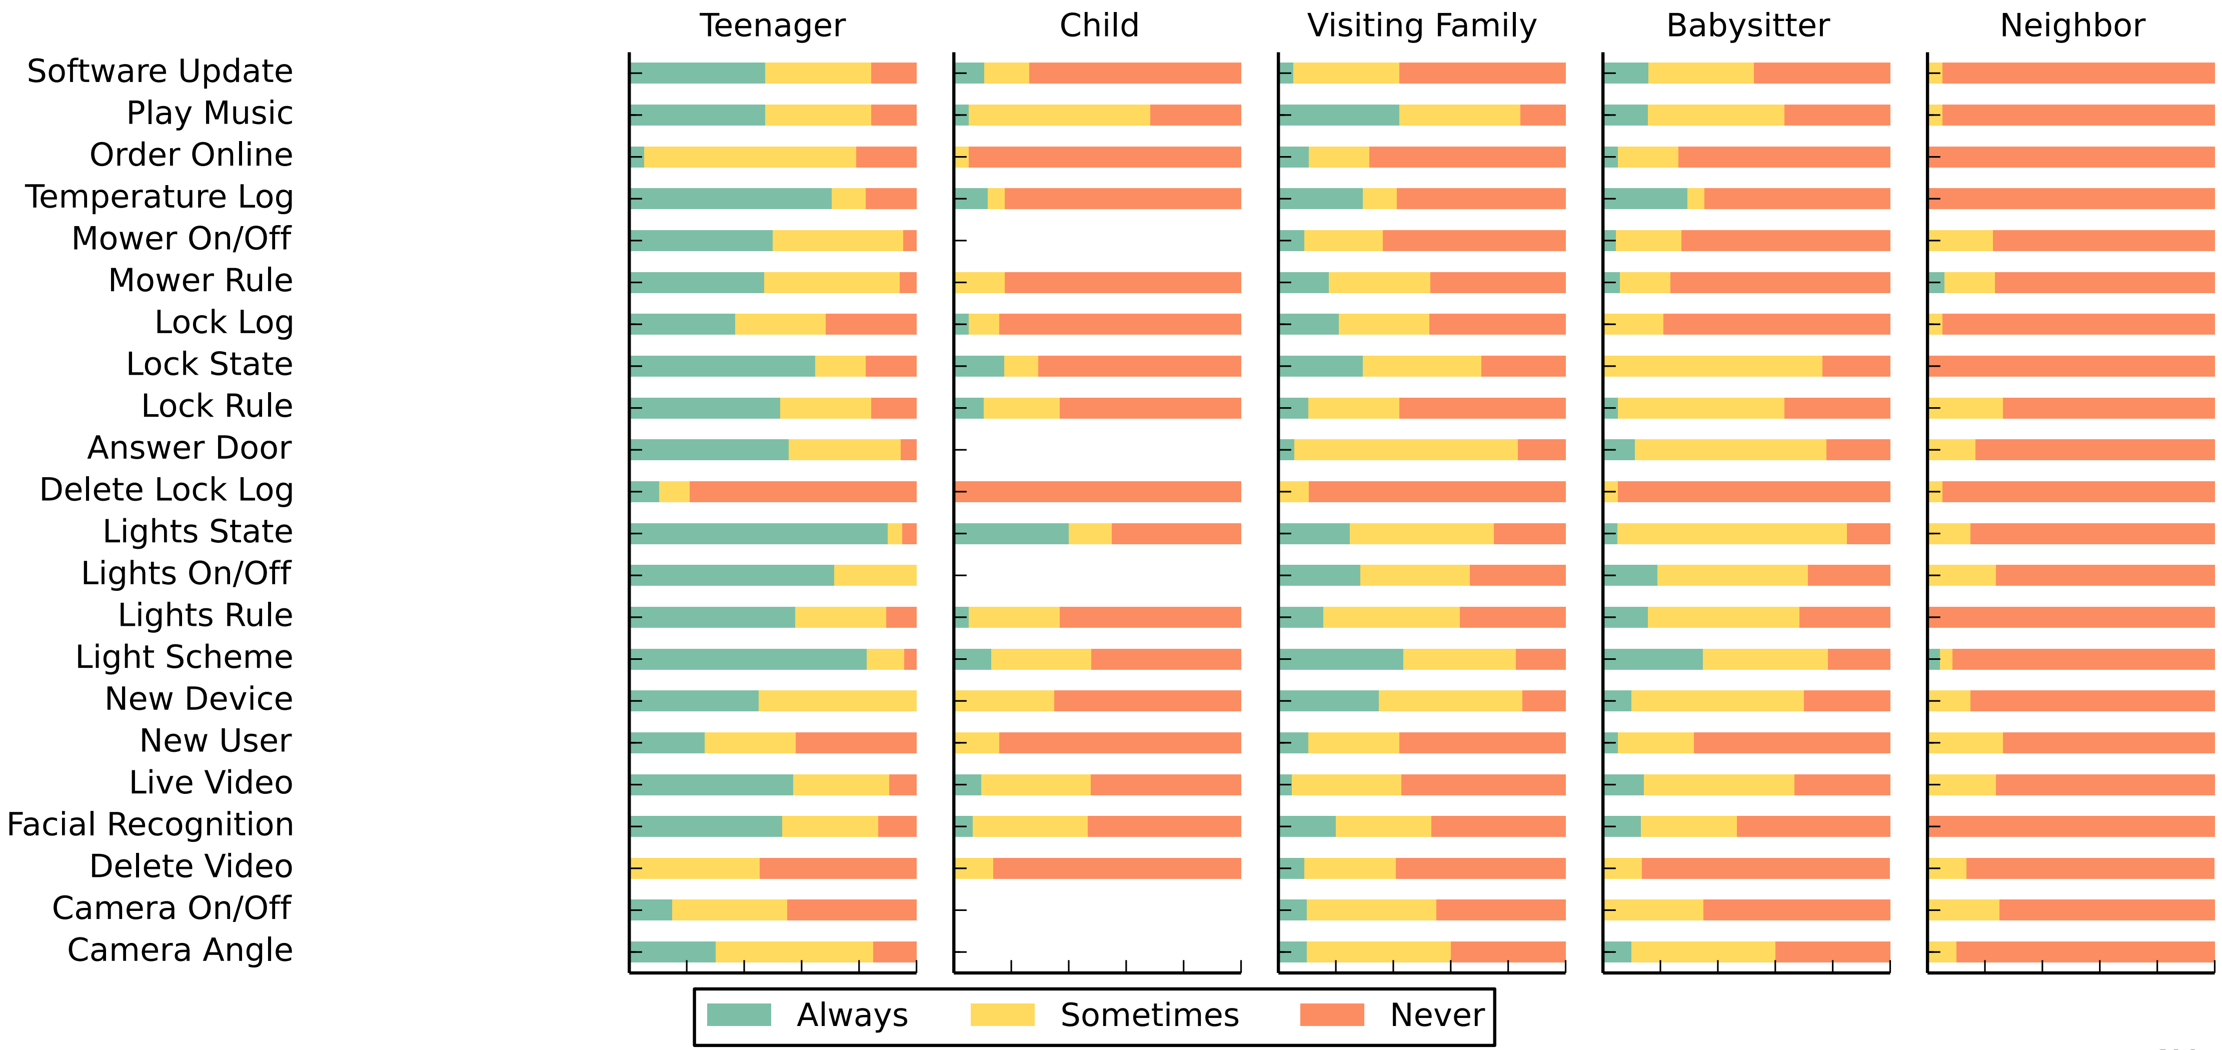

Teenager vs. Child

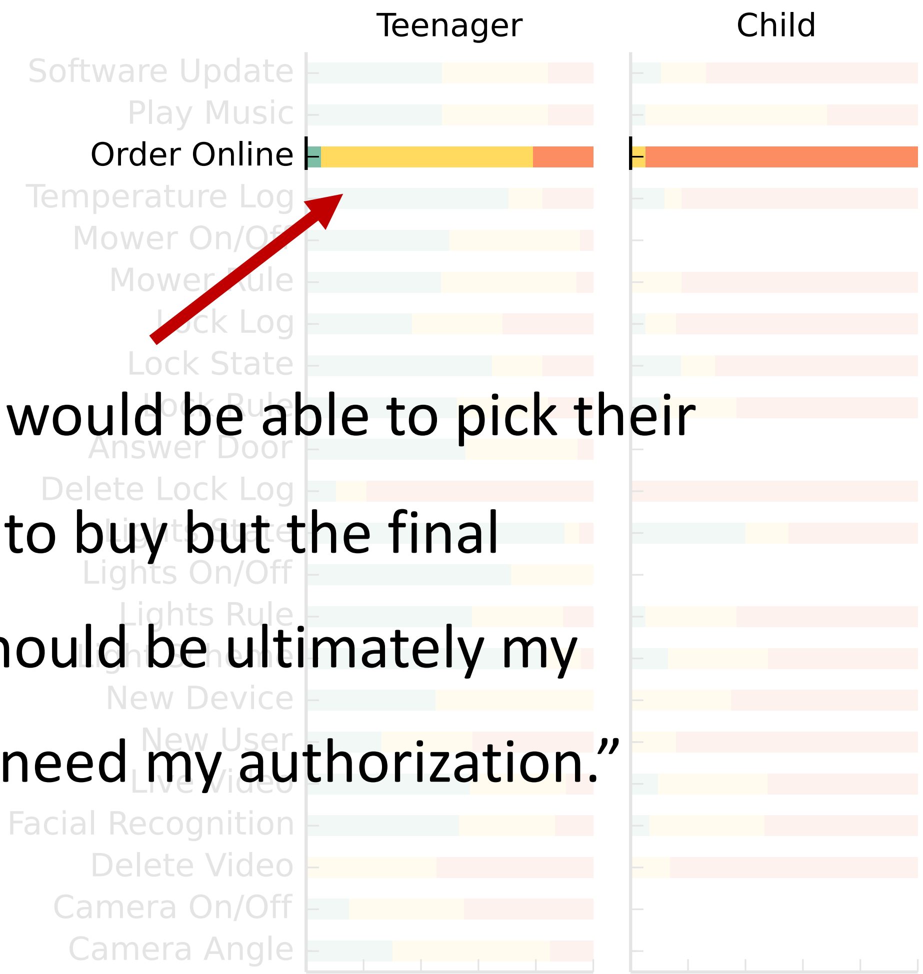

"At 16 they would be able to pick their own things to buy but the final purchase should be ultimately my choice and need my authorization."

Teenager vs. Child

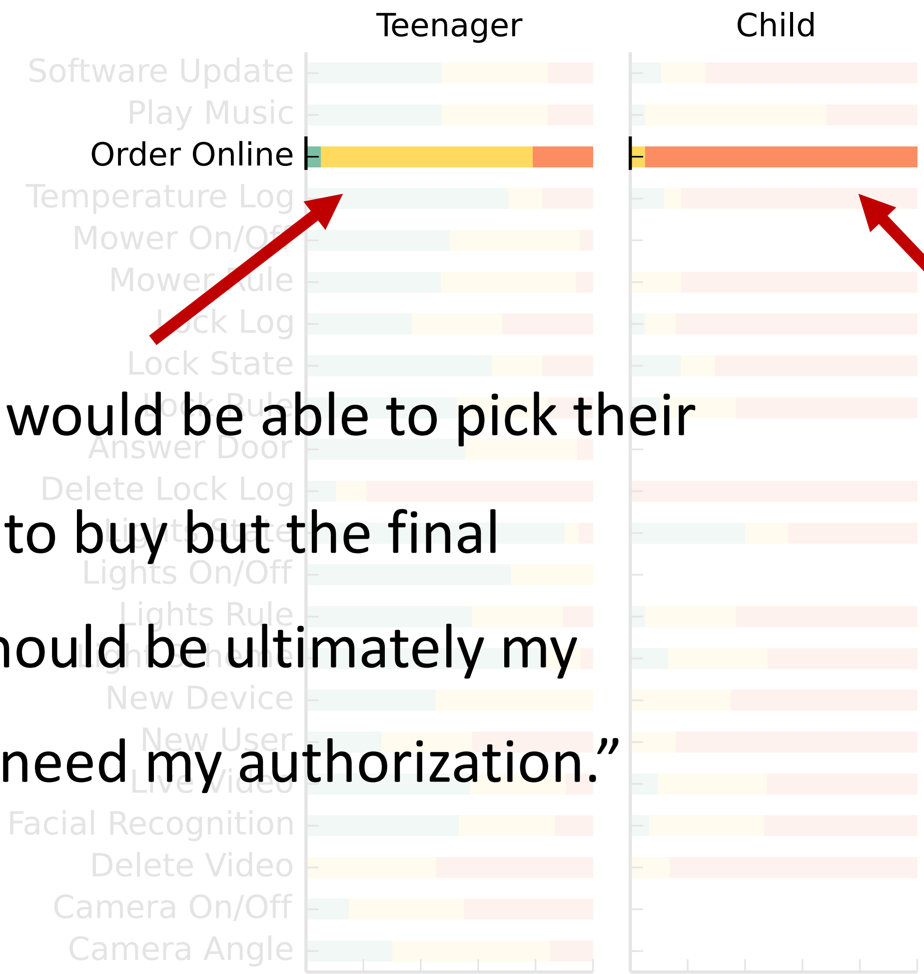

"At 16 they would be able to pick their own things to buy but the final purchase should be ultimately my choice and need my authorization."

"They are in no way responsible enough at this age."

Relationships Matter…But Are Not Enough

Relationships Matter…But Are Not Enough

What does *sometimes* mean?

Contextual Factors

Factor: Time of Day

"I would not want anyone trying to use the mower at night. The neighbors would most likely get mad."

Factor: People Around

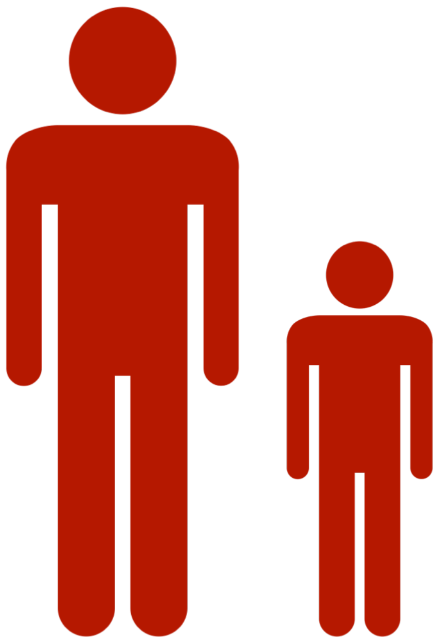

Child

“They would be allowed to use it whenever I am home with them.”

Factor: Location of User

"Why do you need to use it if you aren't close?"

Factor: Location of Device

"If it is used in the bedroom then it would matter who has access."

Factor: Explicit Permission

"When they are authorized by the owner."

Factor: Consequences

Factor: Responsible Usage

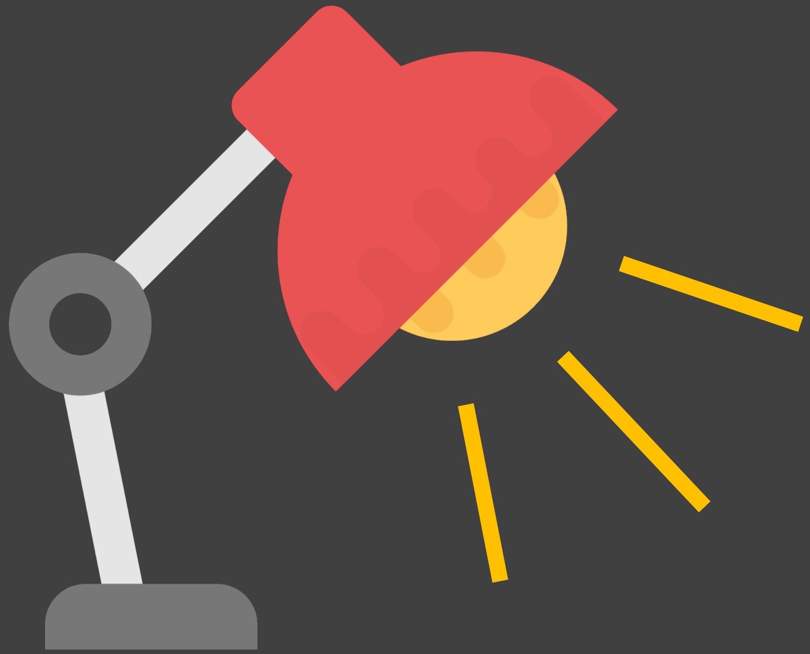

"They shouldn't use the lights if they are using them too frequently."

Icon made by Freepik from www.flaticon.com

Factor: Understanding

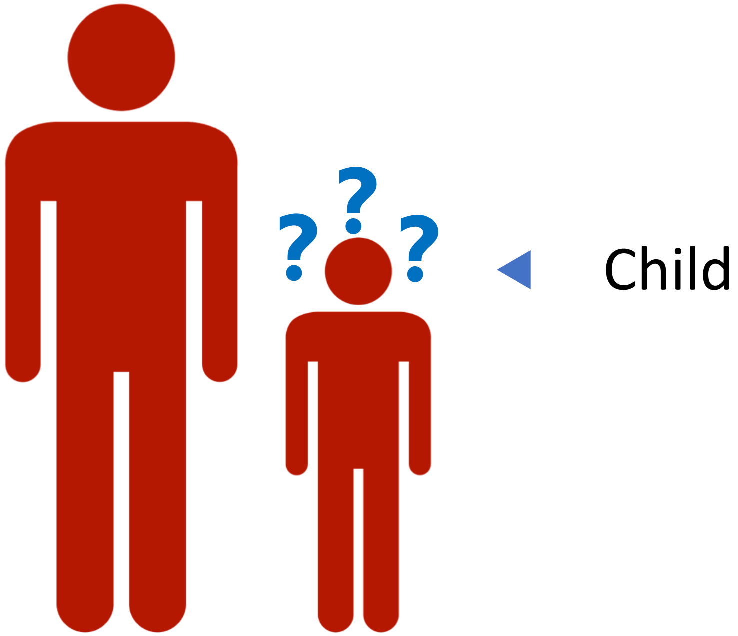

"I would need to teach her how to first."

Factor: Help

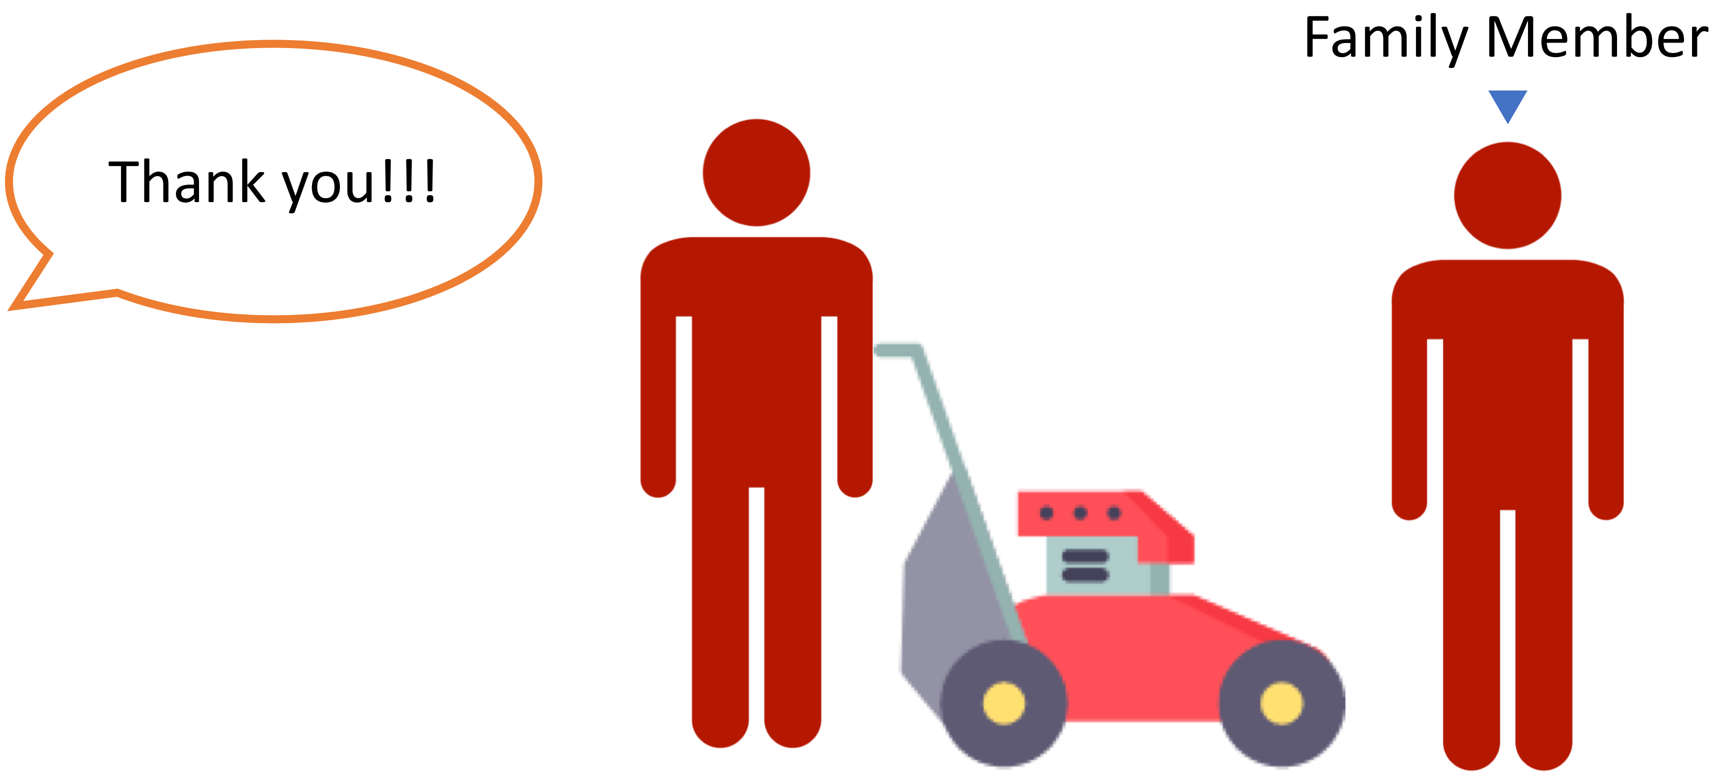

"If they want to come over to mow the lawn, then why not?"

Recap: Missing From Current Systems

Relationships

Capabilities

Contextual Factors

Design Implications

Current: Guest vs. Owner

Smartphone Icon – pixabay.com

Future: Designing for Relationships

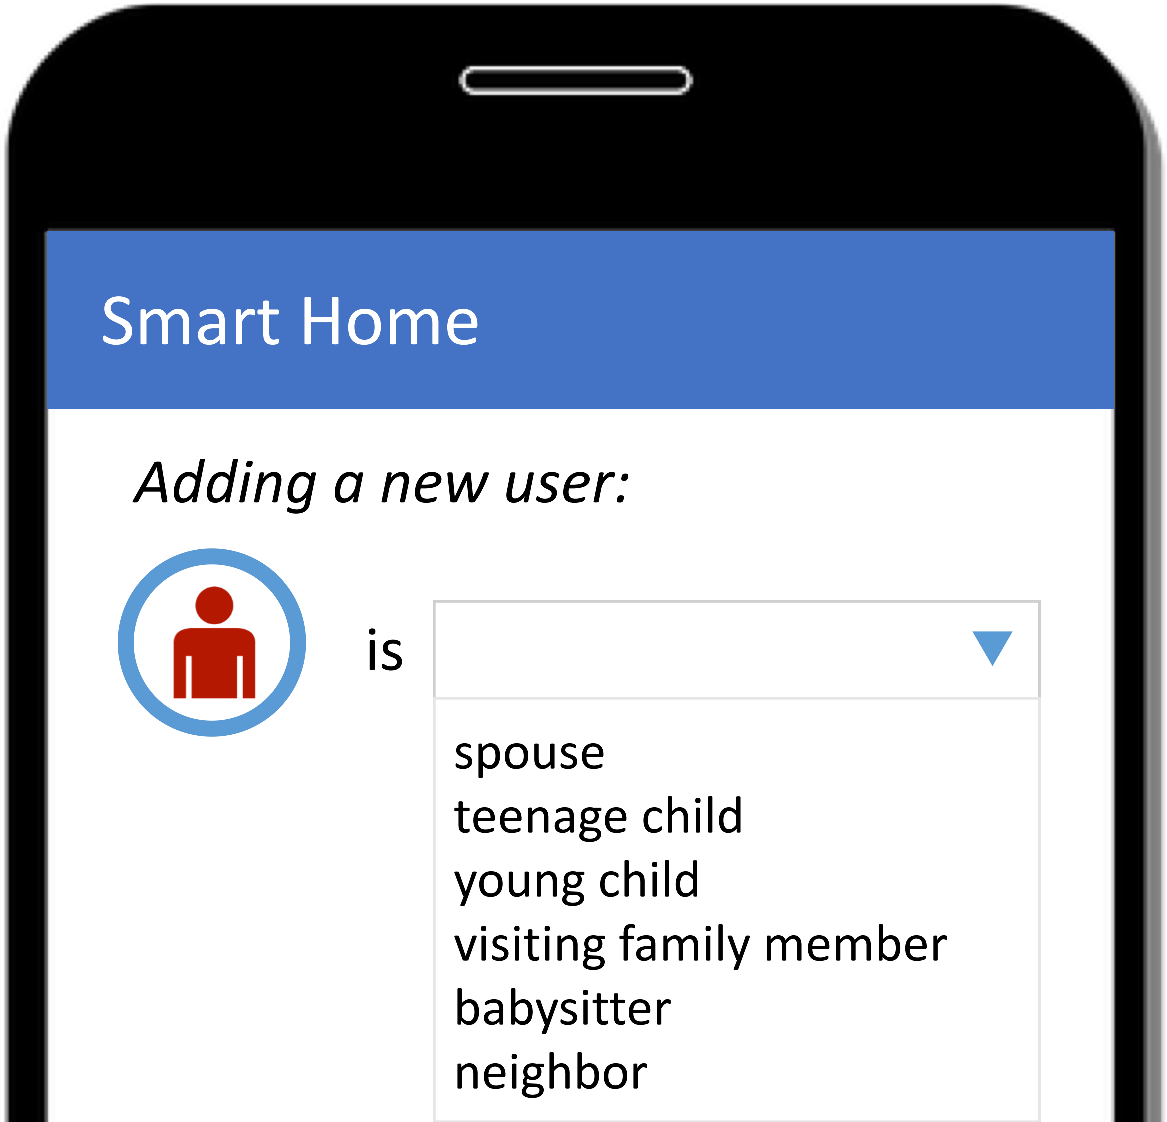

Future: Designing for Relationships

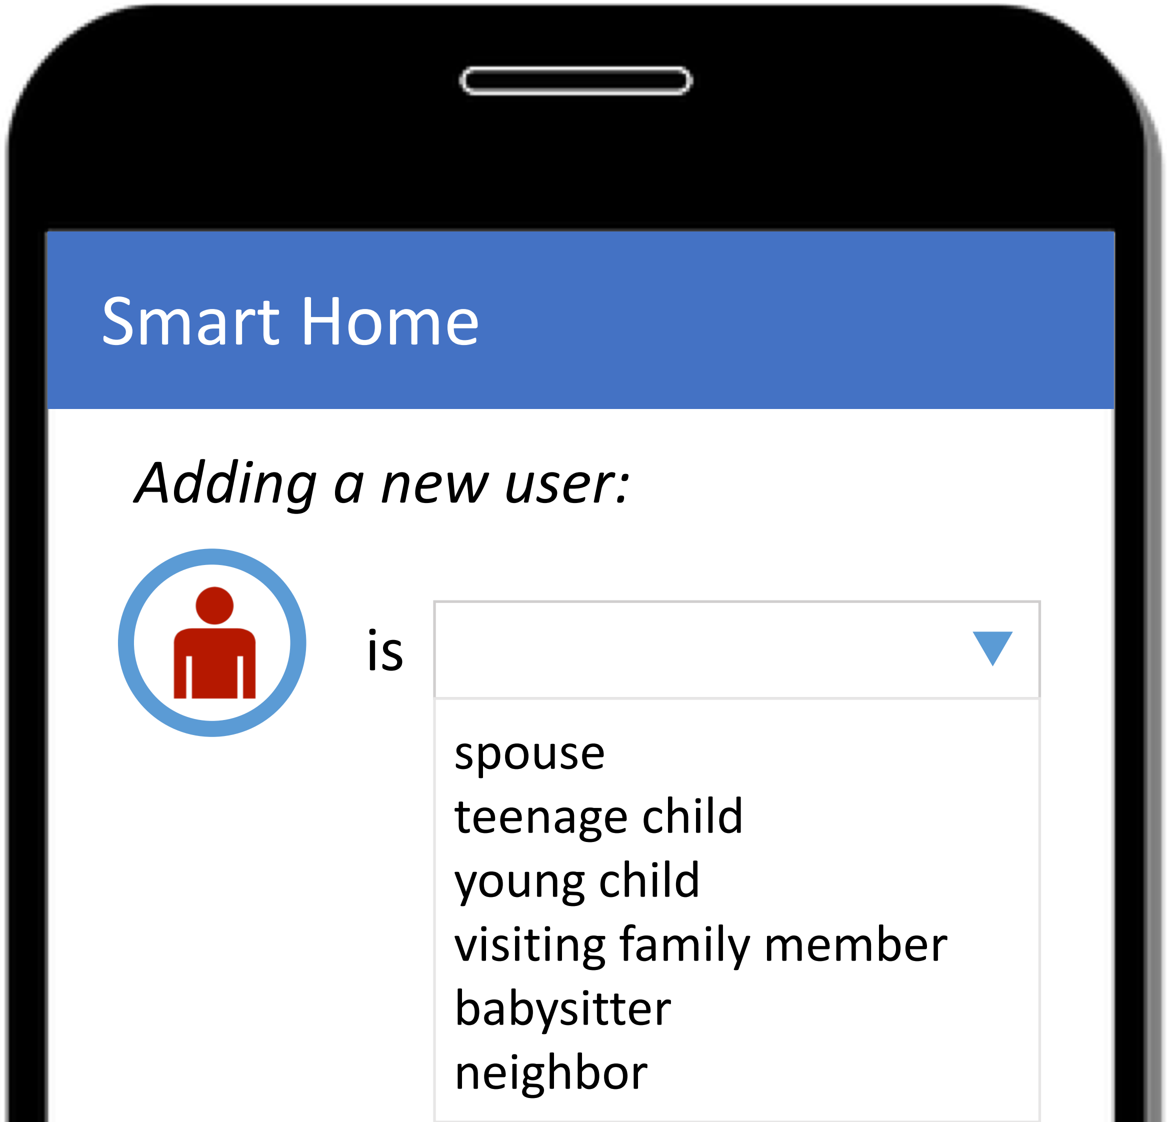

Smartphone Icon – pixabay.com

Future: Designing for Relationships

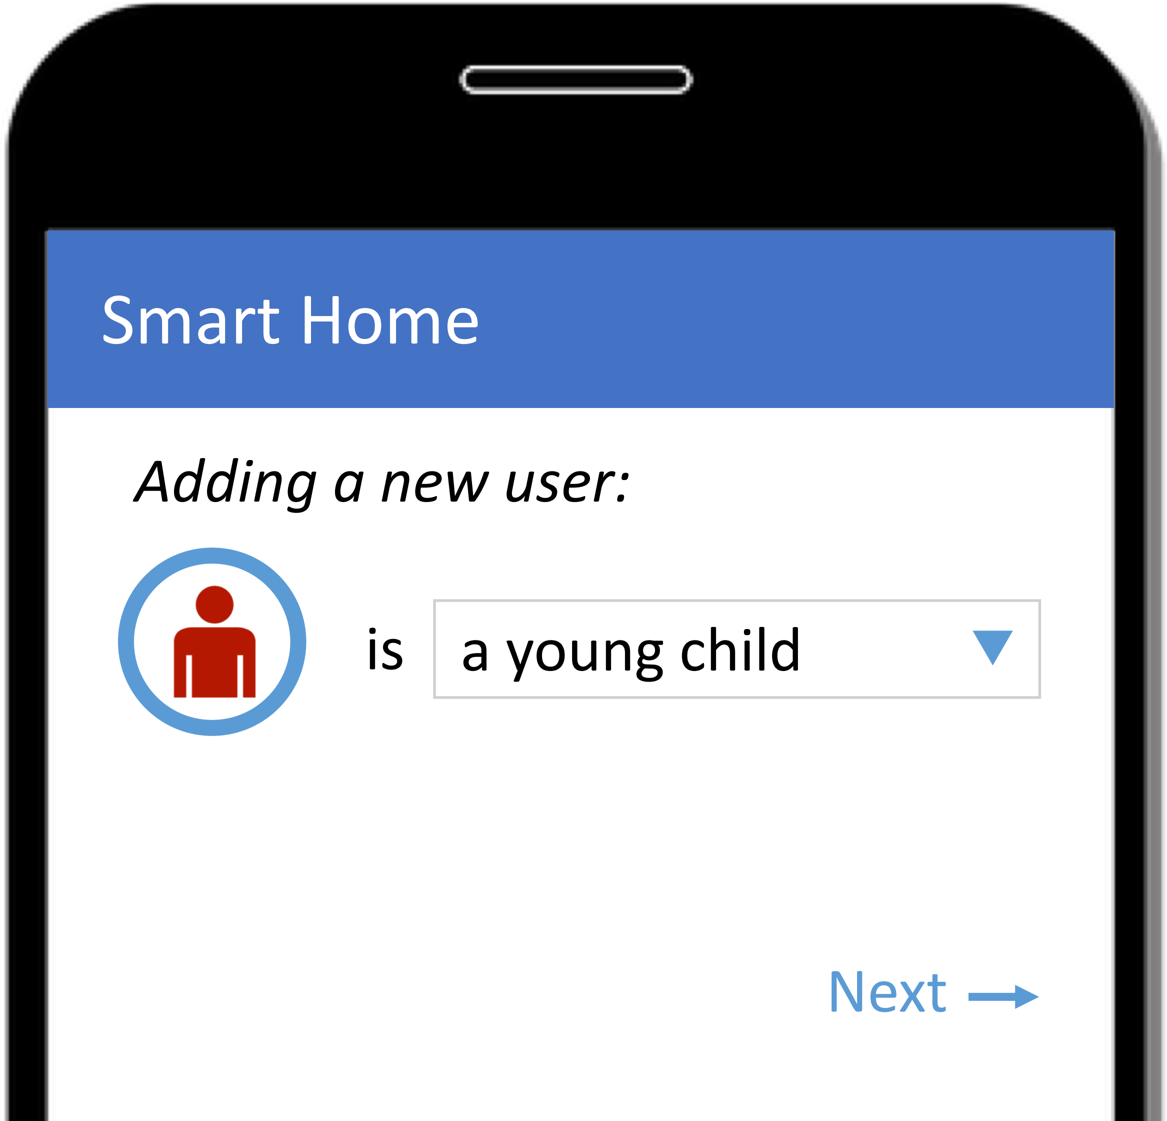

Smartphone Icon – pixabay.com

Future: Relationships and Capabilities

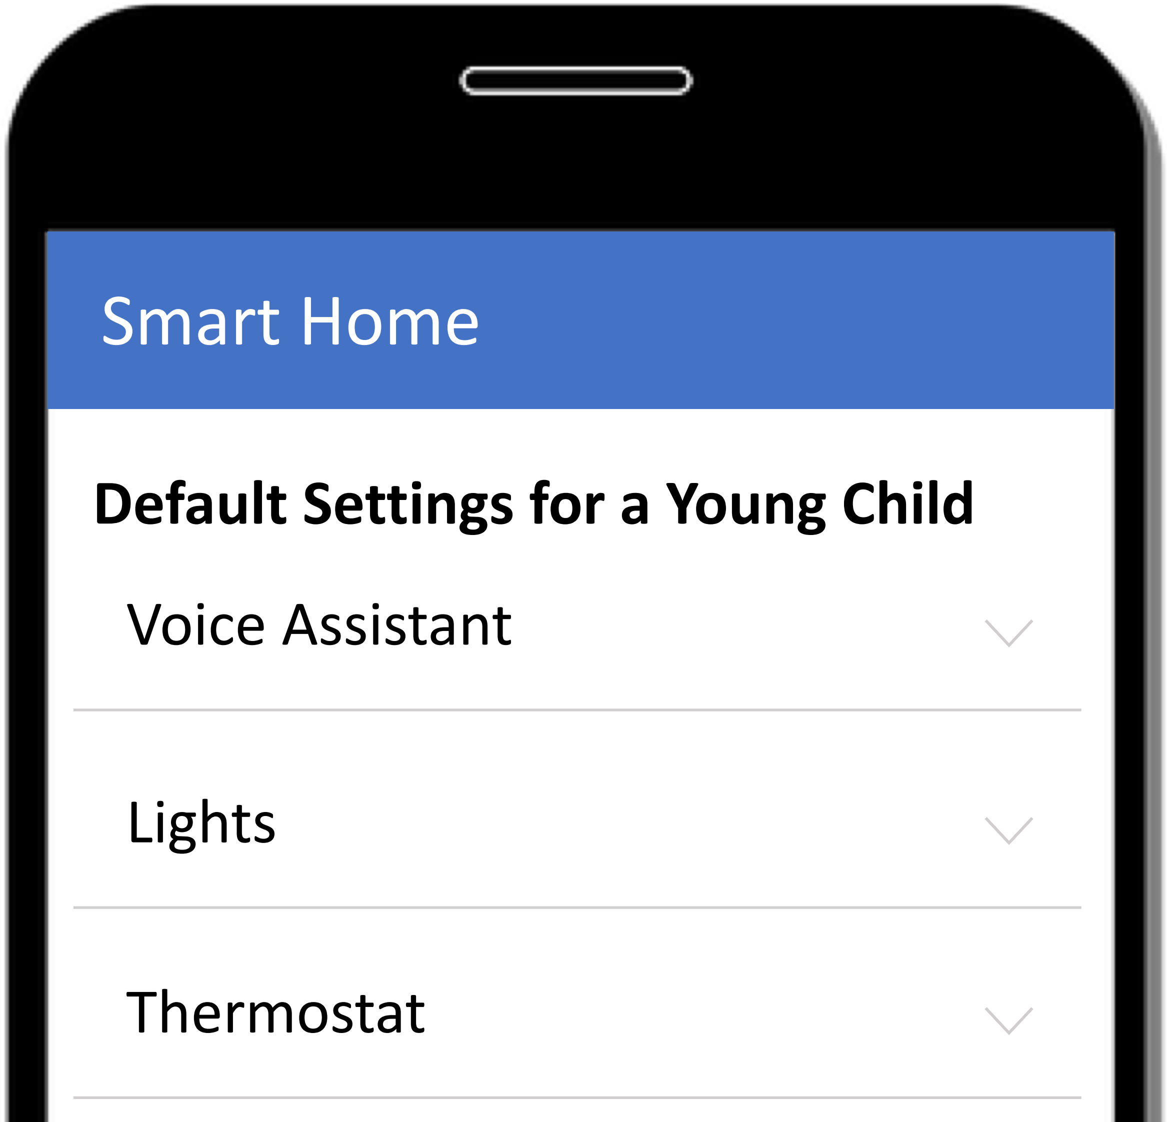

Smartphone Icon – pixabay.com

Future: Relationships and Capabilities

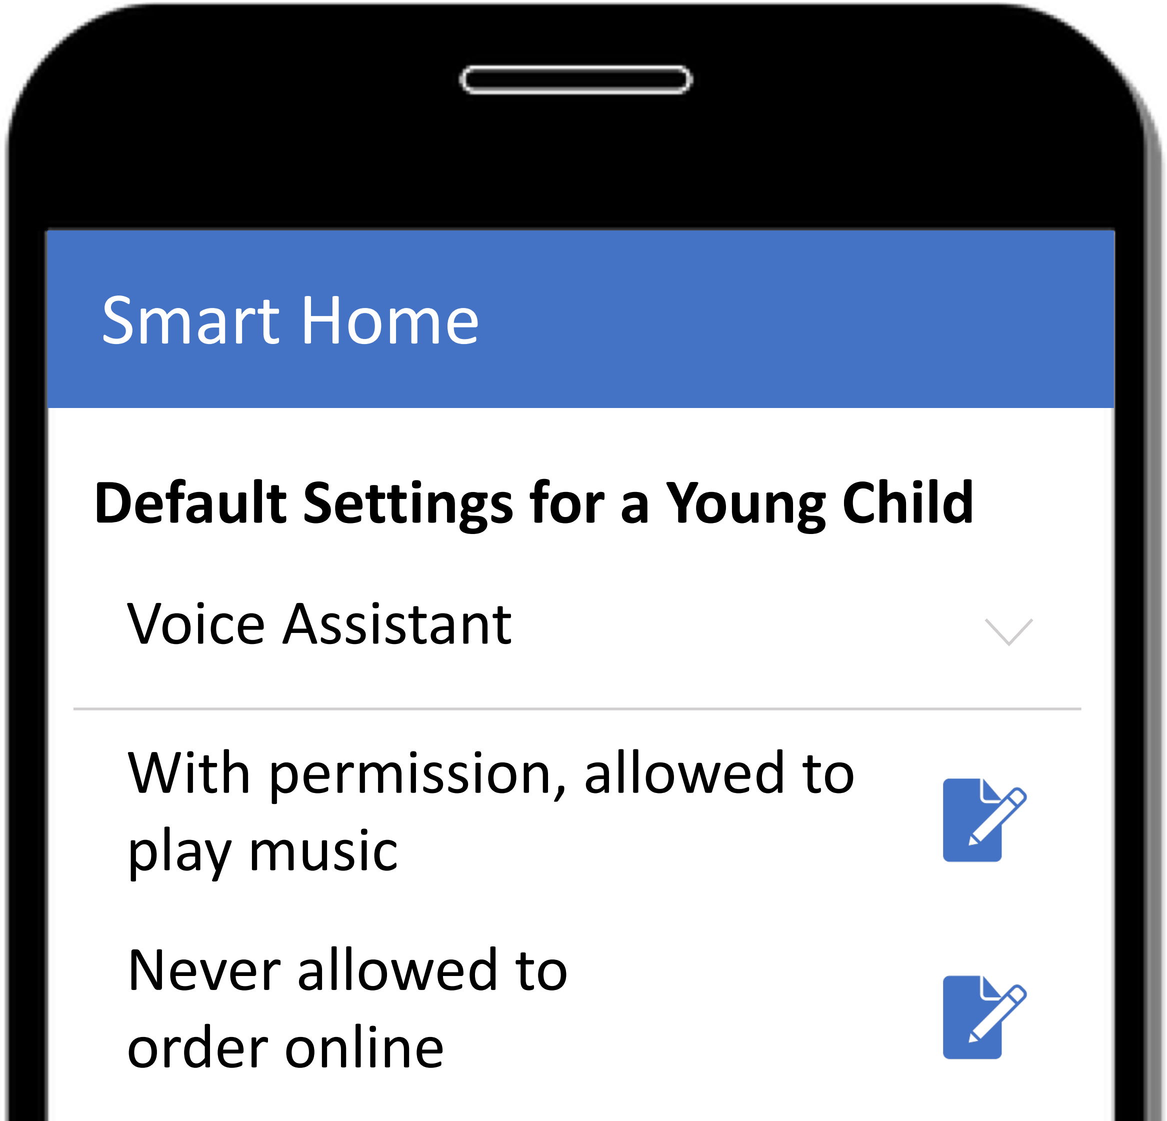

Smartphone Icon – pixabay.com

Current: Full Access or Temporary Access

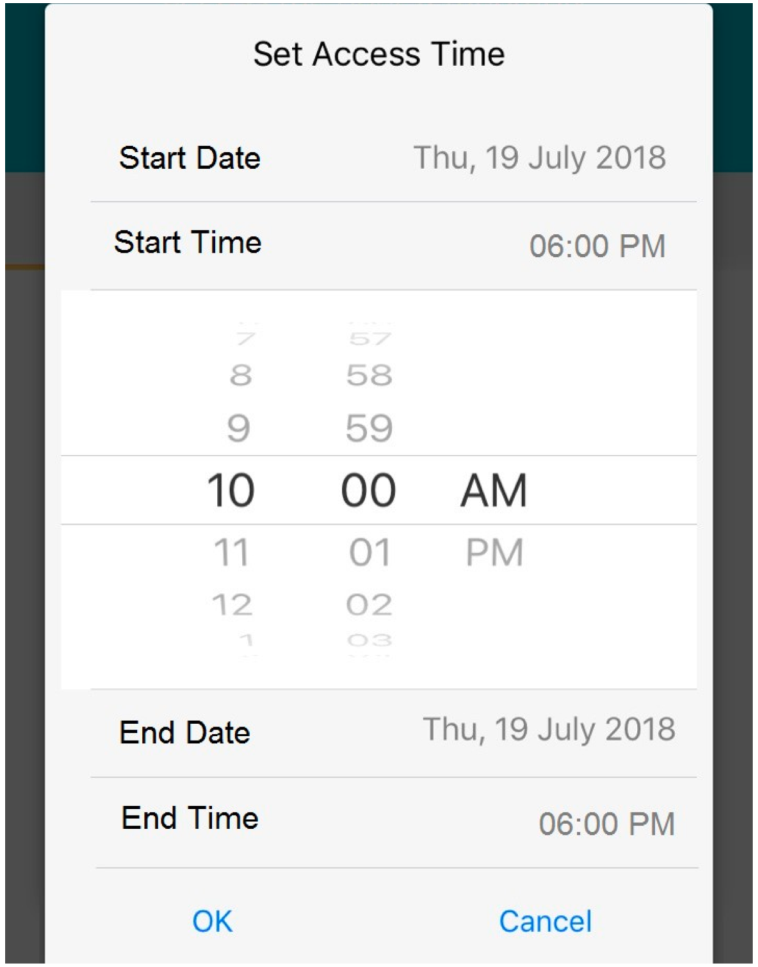

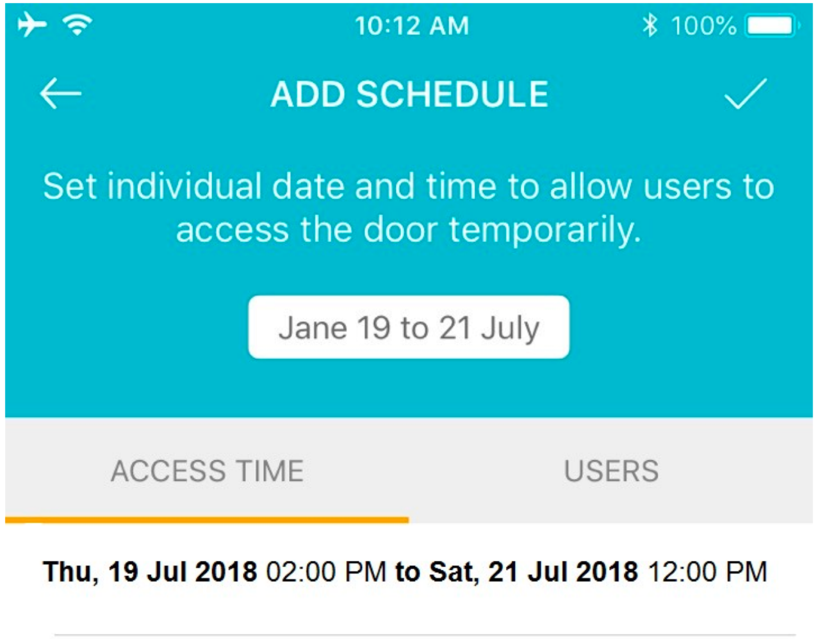

Future: Contextual Factors

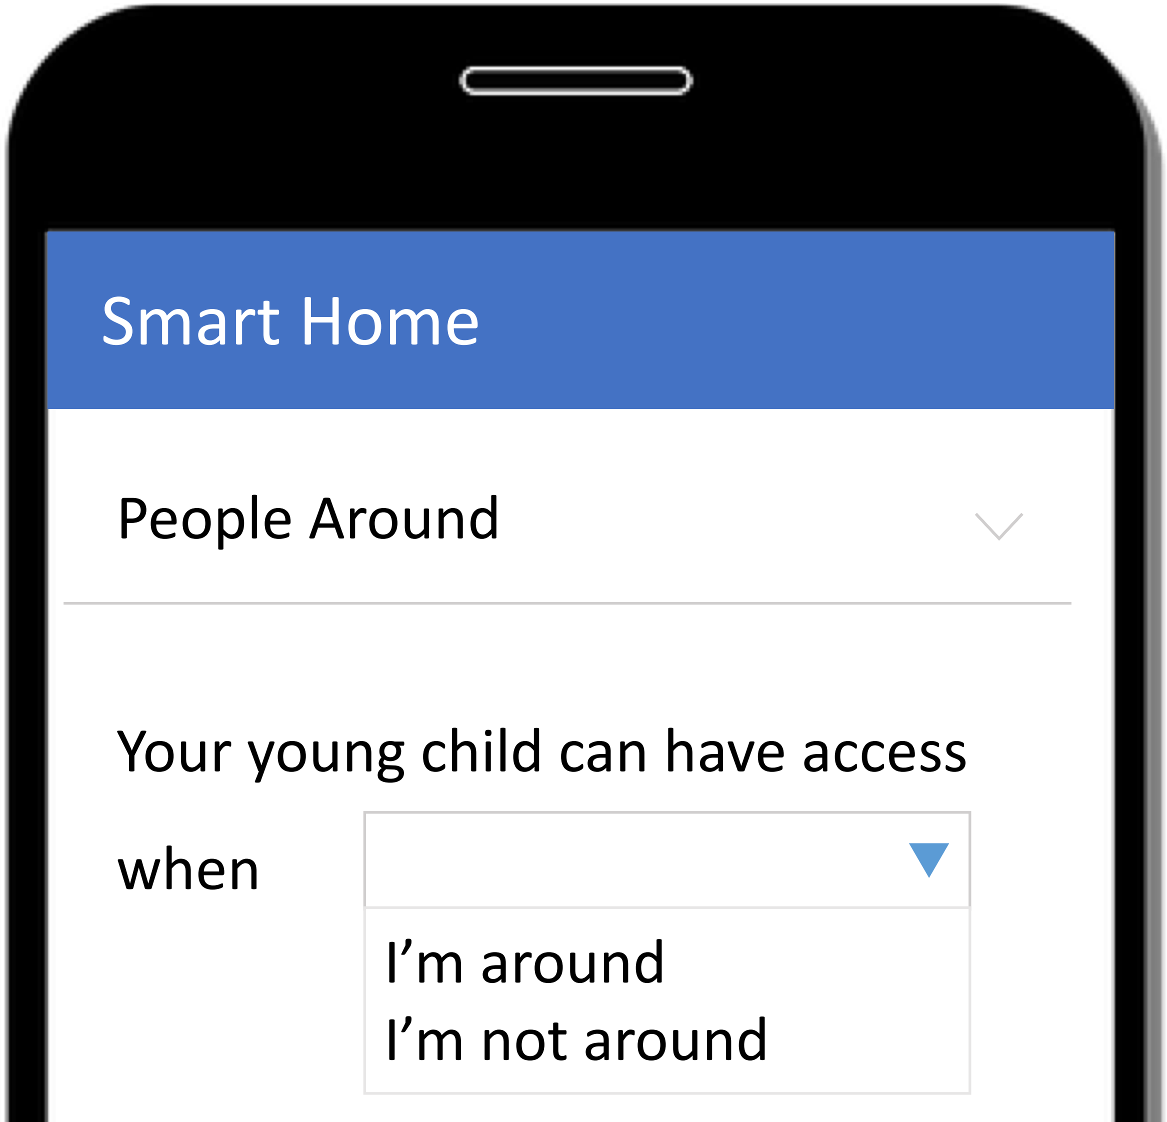

Smartphone Icon – pixabay.com

Future: Device Context

Smart Home

Is your camera an…

Indoor Camera

Outdoor Camera

Smartphone Icon – pixabay.com
Indoor Camera – homesecurity.ml
Blink XT Add On Camera - blinkforhome.com

Future: Device Location

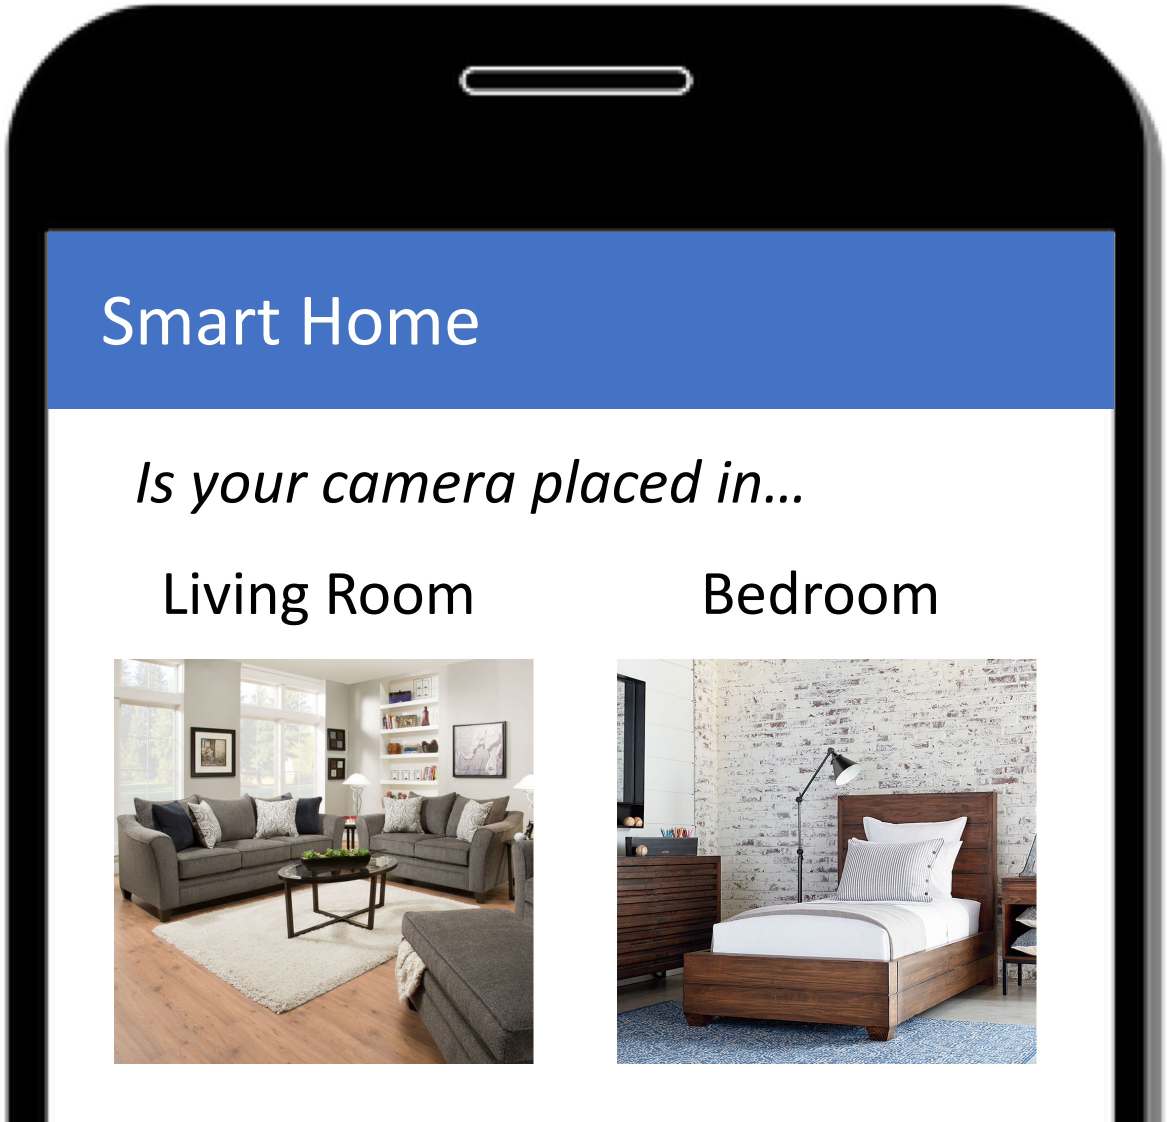

Smartphone Icon – pixabay.com
Living room – raidwarning.net
bedroom-21 - ffooty.com

Capability-Based
Access-Control Policies

Relationships Determine
Default Policies

Support Context-
Dependent Policies

Rethinking Access Control and Authentication for the Home Internet of Things

Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth, Earlence Fernandes, Blase Ur

THE UNIVERSITY OF CHICAGO
RUHR UNIVERSITÄT BOCHUM RUB
UNIVERSITY of WASHINGTON

Fairness and Machine Learning

Galen Harrison
Julia Hanson
Usable Security and Privacy CMSC 23210/33210

1. **Why does this matter?**
2. **What is machine learning?**
3. **Why should we be worried about whether or not it's fair?**
4. **What are some techniques for making machine learning fair?**

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

May 23, 2016

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

Car Insurance Companies Charge Higher Rates in Some Minority Neighborhoods

First-of-its-kind data analysis finds price differences that can't be explained by risk alone

By Julia Angwin, Jeff Larson, Lauren Kirchner, and Surya Mattu of ProPublica
Last updated: April 21, 2017

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

Car Insurance Companies Charge Higher Rates in Some Minority Neighborhoods

First-of-its-kind data analysis finds price differences that can't

Amazon's Face Recognition Falsely Matched 28 Members of Congress With Mugshots

By Jacob Snow, Technology & Civil Liberties Attorney, ACLU of Northern California
JULY 26, 2018 | 8:00 AM

TAGS: Face Recognition Technology, Surveillance Technologies, Privacy & Technology

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

Car Insurance Companies Charge Higher Rates i ds

First-of-it s that can't

The false matches were disproportionately of people of color, including six members of the Congressional Black Caucus, among them civil rights legend Rep. John Lewis (D-Ga.). These results demonstrate why Congress should join the ACLU in calling for a moratorium on law enforcement use of face surveillance.

Amazon's 28 Members of Congress With Mugshots

By Jacob Snow, Technology & Civil Liberties Attorney, ACLU of Northern California
JULY 26, 2018 | 8:00 AM

TAGS: Face Recognition Technology, Surveillance Technologies, Privacy & Technology

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

Amazon scraps secret AI recruiting tool that showed bias against women

Jeffrey Dastin

8 MIN READ

By Jacob Snow, Technology & Civil Liberties Attorney, ACLU of Northern California

JULY 26, 2018 | 8:00 AM

TAGS: Face Recognition Technology, Surveillance Technologies, Privacy & Technology

What is Machine Learning?

- Problem
  - There is some unknown function $f : A \to B$
  - Examples $A = \{\textbf{pictures}\}, B = \{\textbf{is face}\}$

$A = \{\textbf{chess board}\}, B = \{\textbf{optimal move}\}$

What is Machine Learning?

- Can't find f directly, but have examples of $(\overrightarrow{a}, f(\overrightarrow{a}))$
- Can approximate f
- $\overrightarrow{a}$ could be pixels of picture
  - in income prediction (age, education, …)

Linear Regression

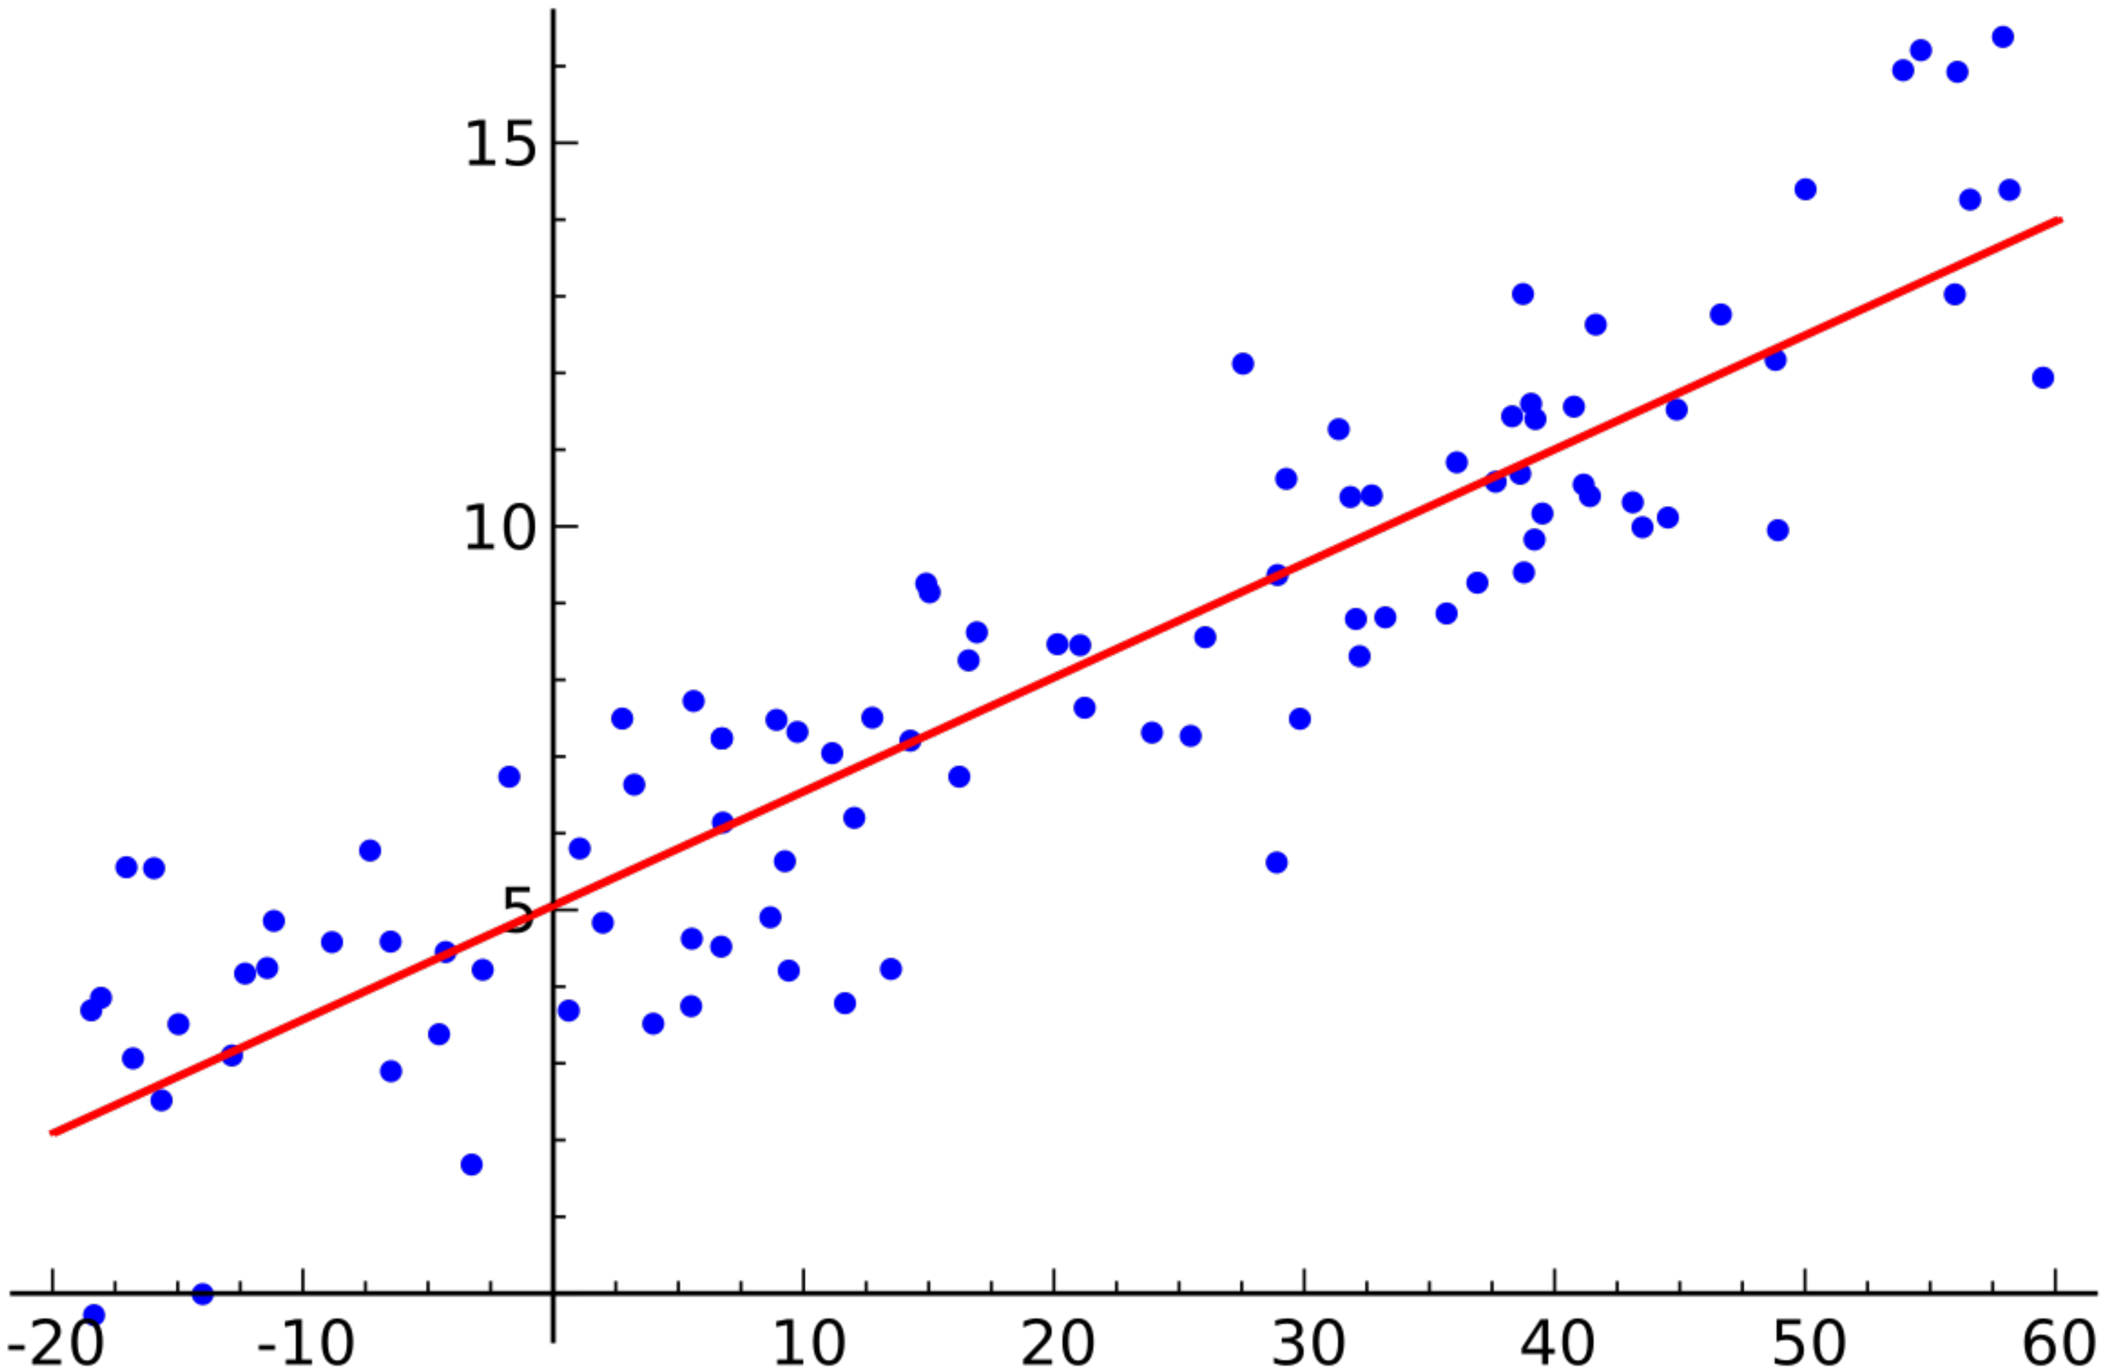

Regression

“Most machine learning is actually regression” - Someone

Key idea: find the right $\overrightarrow{w} = (w_1, w_2, \ldots, w_k)$

Such that $\sum_{i=1}^{n} (w \cdot x_i - y_i)^2$

Other Machine Learning Techniques

- Logistic regression
- Support Vector Machines (SVM)
- Deep learning (aka neural networks)

Key Questions

- These will be more relevant later!
- Does the type of model applied to the problem matter? If so, when?
- When does the machine learning problem matter?
- What, if anything, makes the use of data for ML different from other ways of making decisions?

Returning to Compas

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

May 23, 2016

Propublica reporting on Northpointe risk assessment tool

Risk Assessments

- Risk assessments = Predictive algorithms
- **In bail hearings**: *“will this person commit a crime or fail to appear in court?”*
- **At sentencing:** *“will this person commit crime in the future?”*
- **Theoretical goal: reduce the number of individuals behind bars before trial without increasing risk to the public**

Returning to Compass

137 questions, 10 topics

Current criminal charges	Criminal attitudes
Criminal history	Neighborhood safety
Substance abuse	Criminal history of friends and family
Stability of employment	Quality of social life
Personality	Education and behavior in school

No questions about sensitive features!

Consider the following

$$PPV = \frac{TP}{TP + FP}$$

$$FPR = \frac{p}{1-p}\frac{1-PPV}{PPV}(1-FNR)$$

Chouldechova 2016, Kleinberg et al. 2016

One Problem

- If p differs between two groups, then equal PPV implies differing FPR rates

$$FPR = \frac{p}{1-p}\frac{1-PPV}{PPV}(1-FNR)$$

Other ways bias can arise

- Pre-existing bias
    - Individual - individual people within system design, implementation, use are biased
    - Societal - society as a whole has biases (e.g. a loan system that uses zip codes, reinforcing redlining)

Friedman and Nissenbaum, 1996

Other ways bias can arise

- Technical Bias
  - Computer tools
  - Decontextualized algorithms
  - Random number generation
  - Formalization of human constraints

Friedman and Nissenbaum, 1996

Other ways bias can arise

- Emergent Bias
  - New Societal Knowledge
  - Mismatch between users and system
    - Different expertise
    - Different values

Friedman and Nissenbaum, 1996

Another perspective

(1) *Fair*: lacking biases which create unfair and discriminatory outcomes;
(2) *Accountable*: answerable to the people subject to them;
(3) *Transparent*: open about how, and why, particular decisions were made.

By assuring these conditions are met, we can rest easy, threatened no more by the possibility of an algorithm producing harmful outcomes.

Another perspective

(1) *Fair*: lacking biases … tcomes;
(2) *Accountable*: answer …
(3) *Transparent*: open al … ere made.

By assuring these condi … no more by the possibility of an algorithm producing harn …

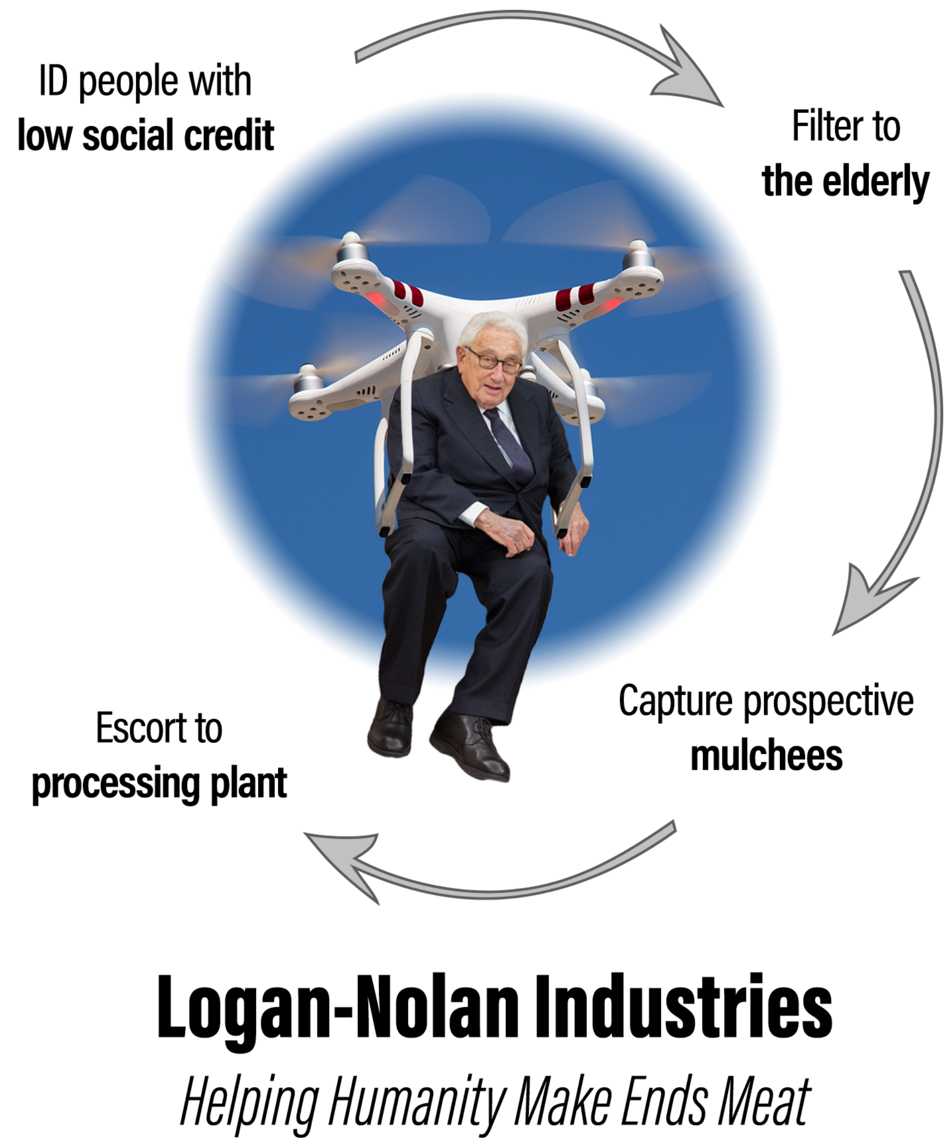

Figure 1: A publicity image for the project, produced by Logan-Nolan Industries

A Mulching Proposal, Keyes et al. 2019

Key Questions

- What is the specific problem that we're trying to solve?
- How much responsibility does the data scientist/machine learning developer have for the broader effects of their work?
- Should we be concerned with *fairness* per se? Or is justice/control/equity a better framing?

Possible solutions

- Technical
- Design
- Regulatory

Individual Fairness

Idea: treat similar people in a similar manner

$M : V \to \Delta(A)$, d_1, d_2 **metrics in** V **and** $\Delta(A)$ **respectively**

$$d_2(M(x), M(y)) \leq d_1(x, y)$$

What intuitions does this encode? What might be some problems?

Fairness Through Awareness, Dwork et al. 2011

Disparate Impact

- Equal Employment Opportunity Commission interprets to say that that if a facially neutral test selects a group at 80% of the rate for other groups, then it is discriminatory according to Title VII of the Civil Rights Act § 2000e-2(a)(2)
- Generalize to $$\frac{Pr(C = 1 \mid X = 0)}{Pr(C = 1 \mid X = 1)} \leq \tau$$

Certifying and Removing Disparate Impact, Feldman et al. 2014

Process Fairness

- Idea: Some features may be fair to use, others may not be
- Base feature use fairness through a survey
- Examples
  - Current charges
  - Criminal History: self
  - Criminal History: social circle
  - Education and school behavior

Human Perceptions of Fairness in Algorithmic Decision Making: A Case Study of Criminal Risk Prediction, Grgić-Hlača et al. 2018

Questions? Comments?

Additional Resources

- ACM Conference on Fairness Accountability and Transparency (ACM FAT*) https://fatconference.org/
- FAT/ML http://www.fatml.org/
- Social Media Collective Critical Algorithm Studies Reading List https://socialmediacollective.org/reading-lists/critical-algorithm-studies/