

Activity Sensing

CS 347

Maneesh Agrawala

Announcements

Discussion sections launched last week

Make sure you are assigned a discussant date

Only one reading for Wednesday

First quiz is in class on Wednesday

During lecture period — bring a pencil

Covers material in lectures and readings through today

Will post example quiz today on Canvas

Email cs347@cs.stanford.edu

Last time

Ubiquitous computing input and output

The typical ubicomp sensing and recognition pipeline

Nontraditional display technologies, augmented reality, virtual reality

HCI interdisciplinarity

Today

Ubicomp envisioning technology in support of our long-term goals

via commodity sensing

via infrastructure-mediated sensing

But what do we do about privacy?

Contextual integrity

HCI 101: tasks and usability

The traditional frame of human-centered design has focused on improving **usability** for well-defined **tasks**, especially tasks of short duration and focused attention

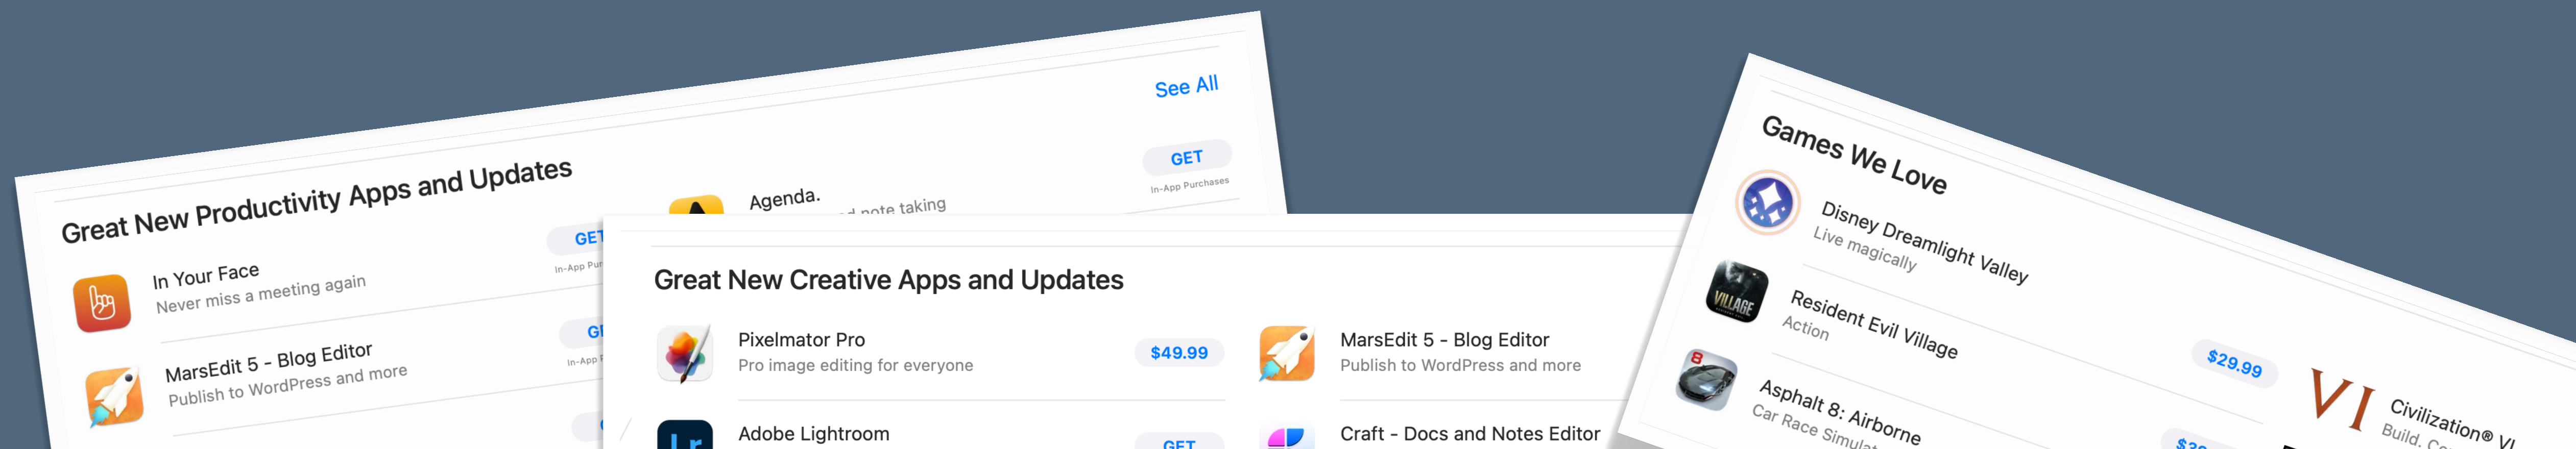
Result: HCI, UI design, usability are now commonplace in industry and academia



HCI 101: tasks and usability

But, this success has come **at a cost**: a focus on interaction design and usability that is not backgrounded as per ubiquitous computing, but **apps at the forefront** of our attention

This is a **legitimacy trap**: what we used to argue for the importance and legitimacy of HCI—task-based usability—is now holding us back [Dourish 2019]



Ubiquitous computing's response: “Hell, no”

This reductive view of HCI as app-ification is limiting

As technology diffuses into all aspects of our lives, its biggest impact may be not on short-lived tasks on screens, but in **issues of much greater societal importance**—education, health, sustainability—and **issues facing a wider cross-section of the population**

This position entails a lens on what a design might encompass—what is, and isn't in scope—far beyond typical app bounds

Behavior change and HCI

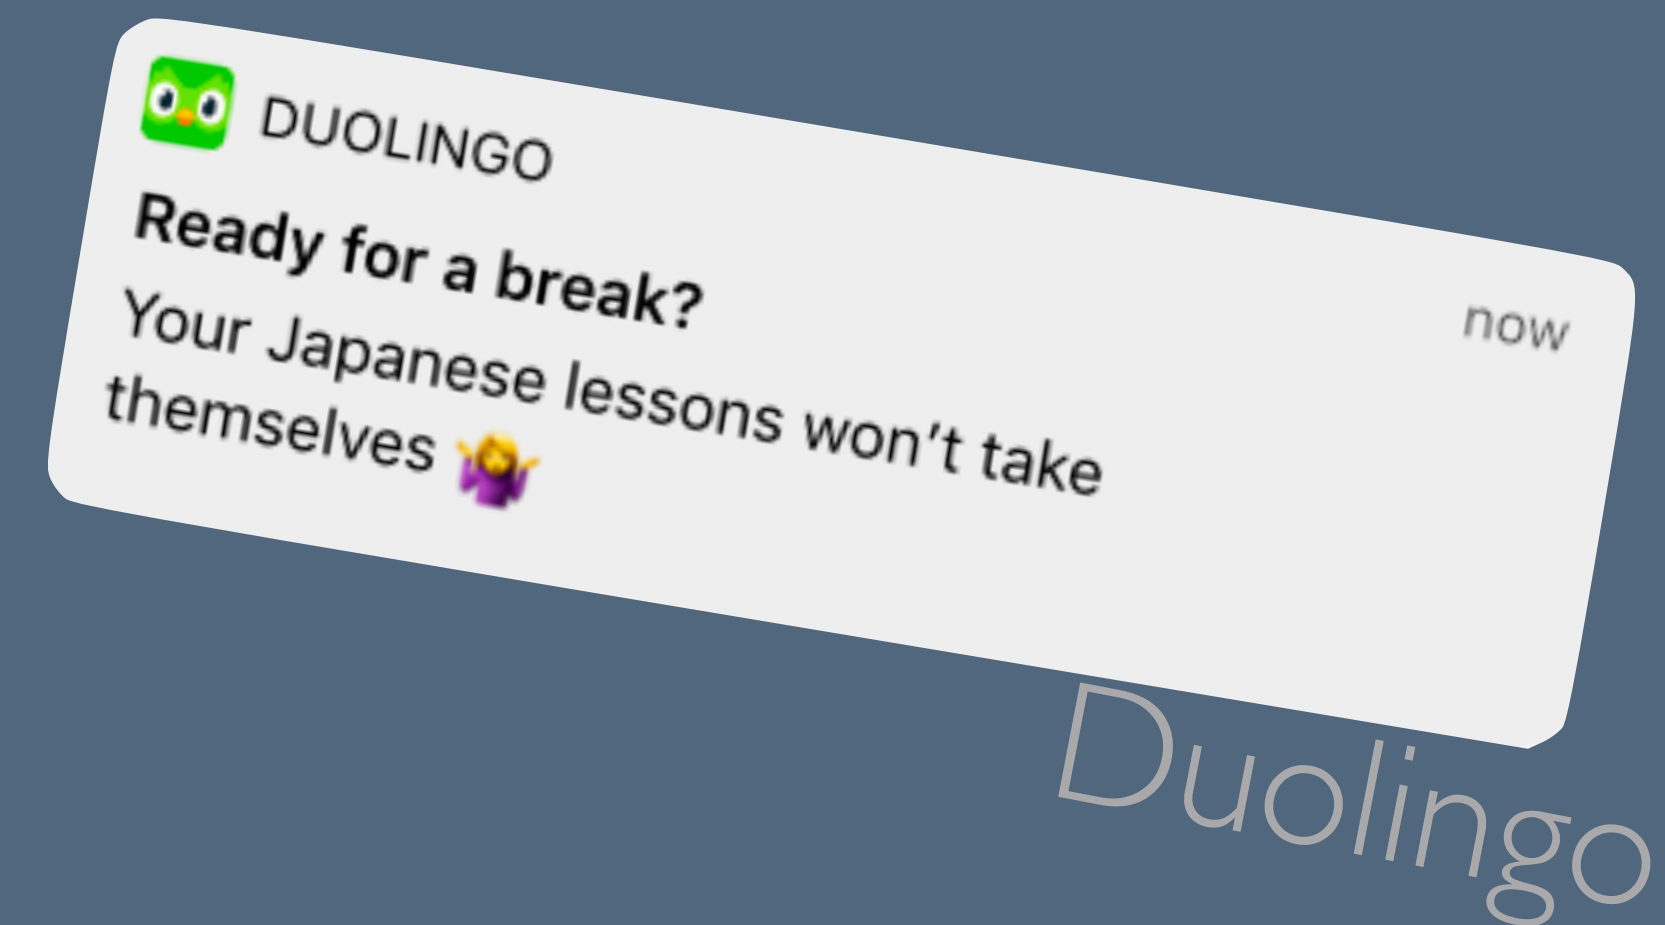
Many of the goals in today's lecture fit under the heading of **behavior change**: designs that shape what we do and when

Change in behavior usually requires new interface design:

“If I only knew how much I was doing, I'd stop.”

“If only I got a reminder at the right time...”

“If only I could know if my elderly grandmother were at risk...”



Where to focus?

Identify a long-lived activity or resilient societal challenge

Ask whether computing can help produce the data or intervention necessary to move the needle in a meaningful way

The answer may be “no”!

But sometimes it's

“Yes!”



Commodity sensing:
repurposing existing
hardware

Typical recipe: unobtrusive and commodity

“Can we **unobtrusively detect** _____ using **commodity smartphones**?”


Unobtrusive: without much active user participation

Commodity: widely available and mass produced

Using a similar recognition pipeline as the previous input lecture

...and potentially using novel sensors that could feasibly be integrated into a smartphone or smartwatch in the future

Physical health



Can we monitor blood pressure using commodity smartphones? [Wang et al. 2018a]

Yes: measure the time between the heart pumping (via phone accelerometer) and the blood moving in an artery in your finger (via phone camera with flashlight on)

Can we detect opioid overdose — breathing cessation — with commodity smartphones? [Nandakumar, Goldakota, Sunshine 2019]

Yes: emit an inaudible frequency sweep (FMCW). It bounces off the person and returns to the phone's mic. The chest moving in and out modulates the time to return, from which we derive a breathing rate

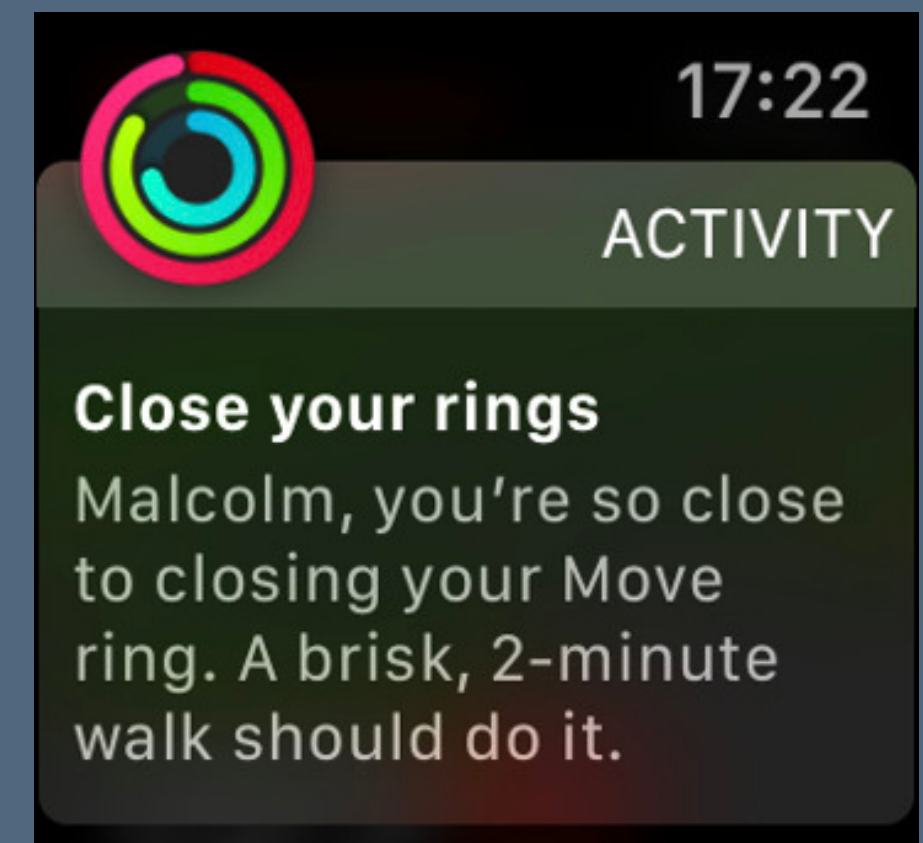
Physical health

Ubifit [Consolvo et al. 2008]: the first system to show that exercise interventions could work with commodity sensors and readily-available glanceable interfaces over long periods

Not yet deployed in industry, but frontiers:

Detect when someone is eating, using inertial measurements on a smartwatch [Thomaz, Essa, Abowd 2015]

Data-driven answers to: Do I gain weight when I have busy days? Do I walk more when I work in the city? Do I sleep better on nights after I work out? [Bentley et al. 2013]



Apple Watch
[via AppleInsider]



Mental health

Stress: can we detect stress levels by listening to your voice [Lu et al. 2012], or by how tightly you're gripping the mouse+keyboard at a computer [Hernandez et al. 2014] or your steering wheel in your car? [Paredes et al. 2018]

Depression: Can we detect depression symptoms using commodity smartphones? [Wang et al. 2018b, Xu et al. 2019] Loneliness and social isolation? [Doryab et al. 2019]

Mental changes: Can we detect mental health changes such as psychotic relapse before they're typically diagnosed? [Ben-Zeev 2017; Wang et al. 2016]

The answer to all of these questions is now a (qualified) yes.

Elder care

How might we design technologies to support successful aging in place? [Kidd et al. 1999]

Can we detect...

Falls, without smartwatches? [Palipana et al. 2018]

Levels of movement and activity in the home?

Neurodiversity

Record and track care for people conditions such as autism

Interventions might include:

Reducing the effort for capturing data about children with autism [Kientz et al. 2007]

Creating interactive tools to aid communication with caretakers [Hayes et al. 2010]

Tools for practicing social skills [Escobedo et al. 2012]

**Infrastructure-mediated
sensing and societal goals**

Typical recipe, part deux

“Can we **unobtrusively detect** _____ using **a sensor in the environment?**”

Goal: avoid needing to instrument people in any way (unobtrusive)

Again using a machine learning classification pipeline

Typically, we achieve this by leveraging infrastructure already available in the environment. This is referred to as **infrastructure-mediated sensing** [Patel et al. 2008]

Sustainability

One major challenge is knowing where my energy and power is going: **which appliances and activities are driving most of my consumption?**

If we knew that, we could identify ways to reduce our energy or water use.

But we can't go around fitting every socket with a sensor...

Disaggregating electrical use

[Patel et al. 2007]

Can we track appliance usage without complex installation or many invasive sensors?


Plug a sensor into a single plug in your home and listen to electrical noise on the power line when switched or in operation

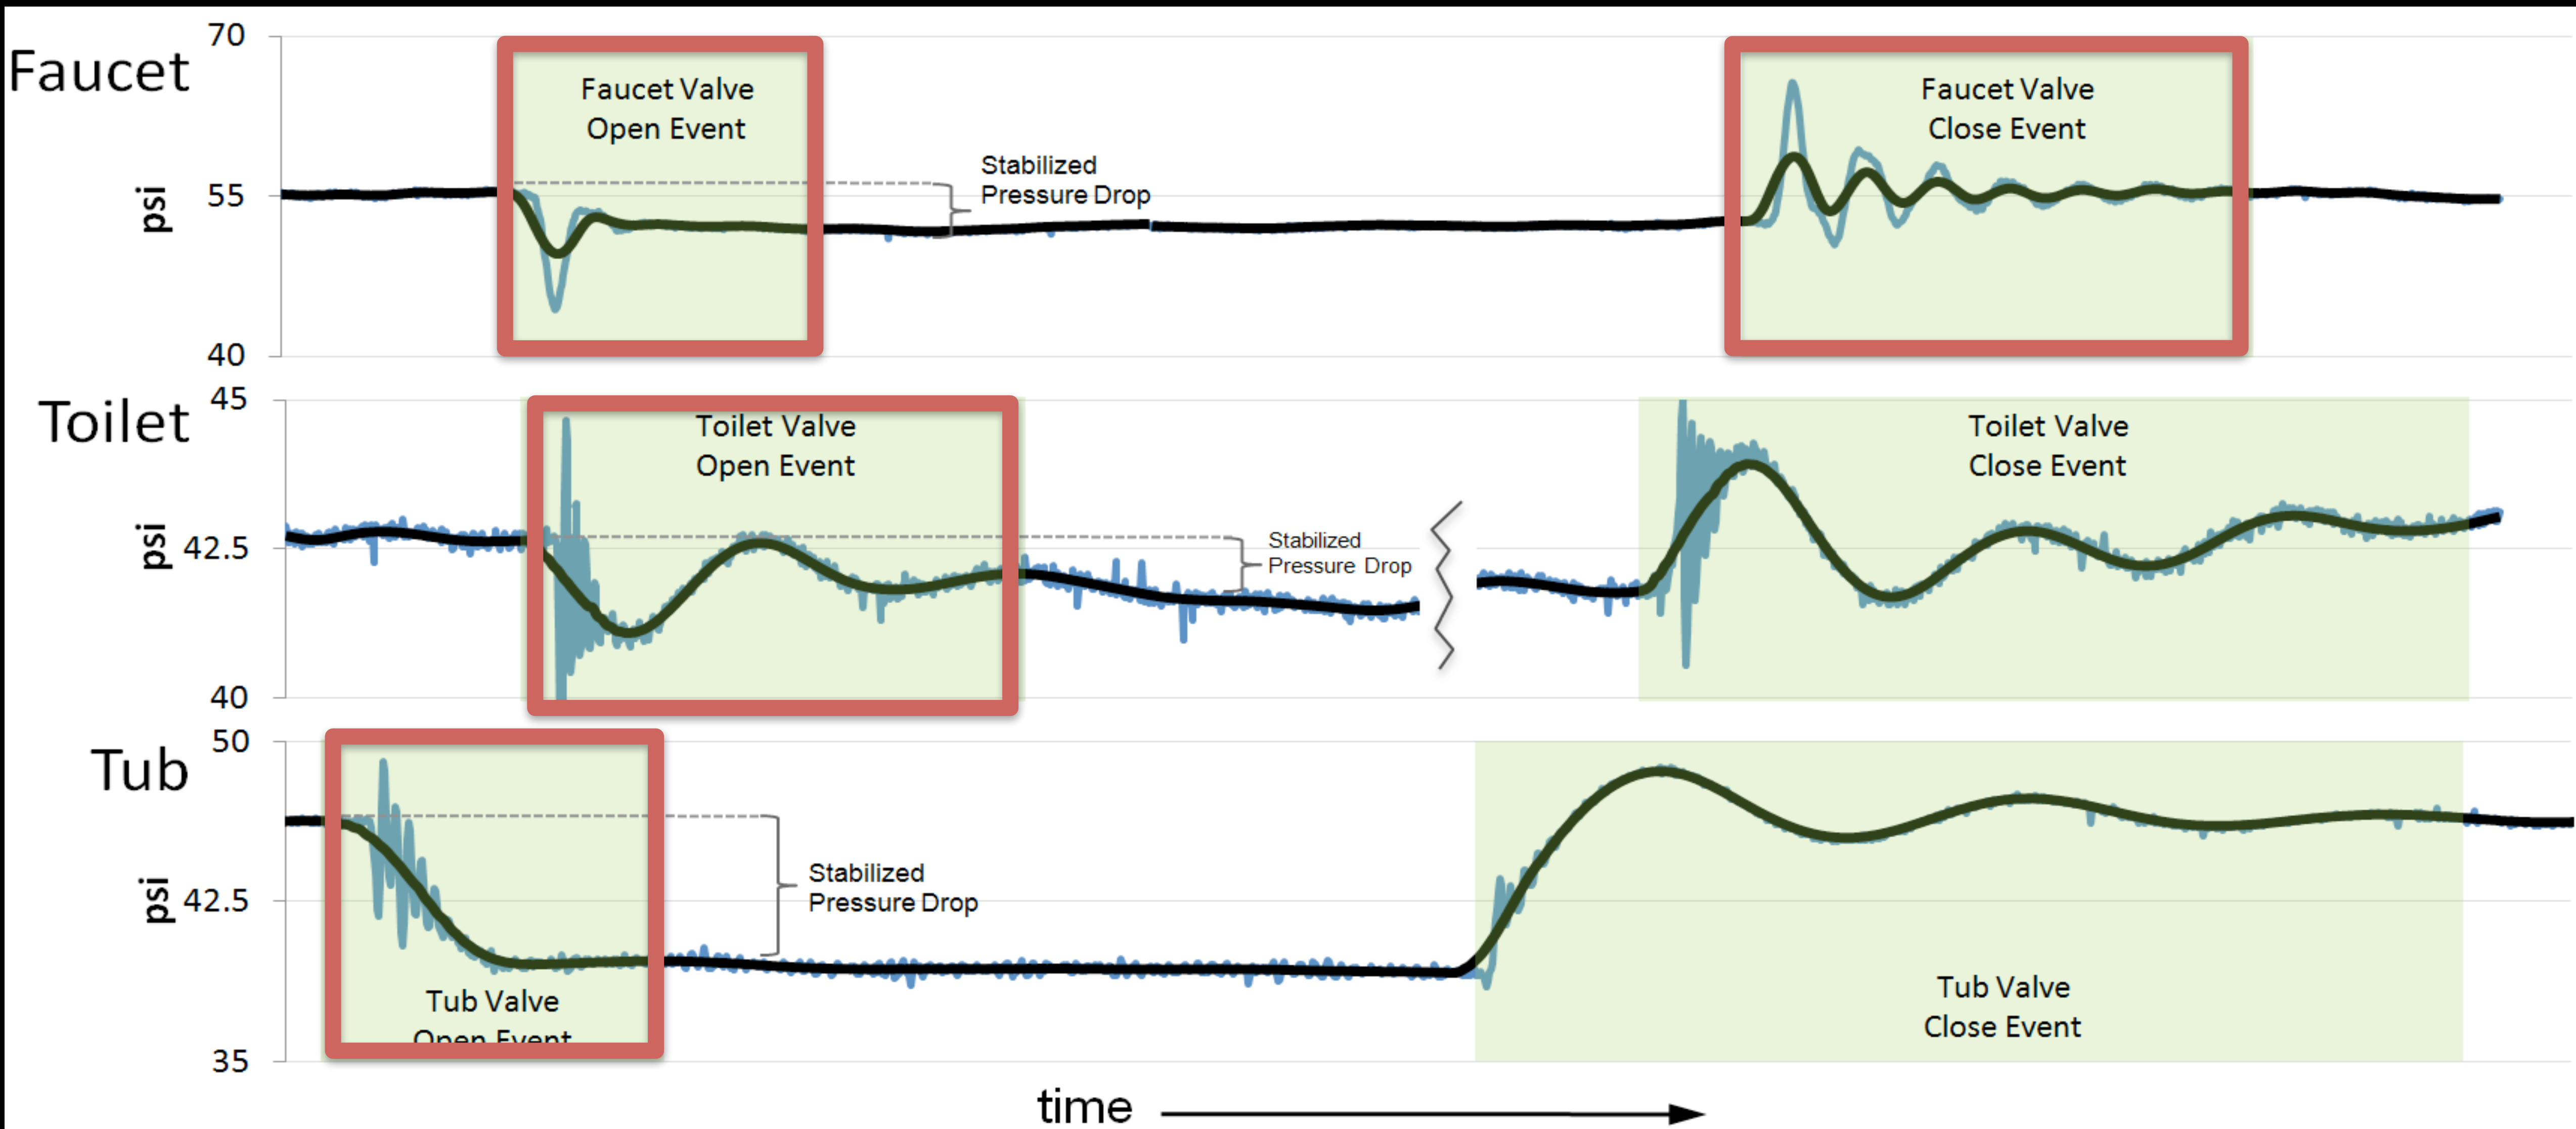


Disaggregating water use

[Froehlich et al. 2009]

A single pressure sensor attached to a hose outdoor faucet
Since your water pipes are typically all connected, that one sensor can see a lot...



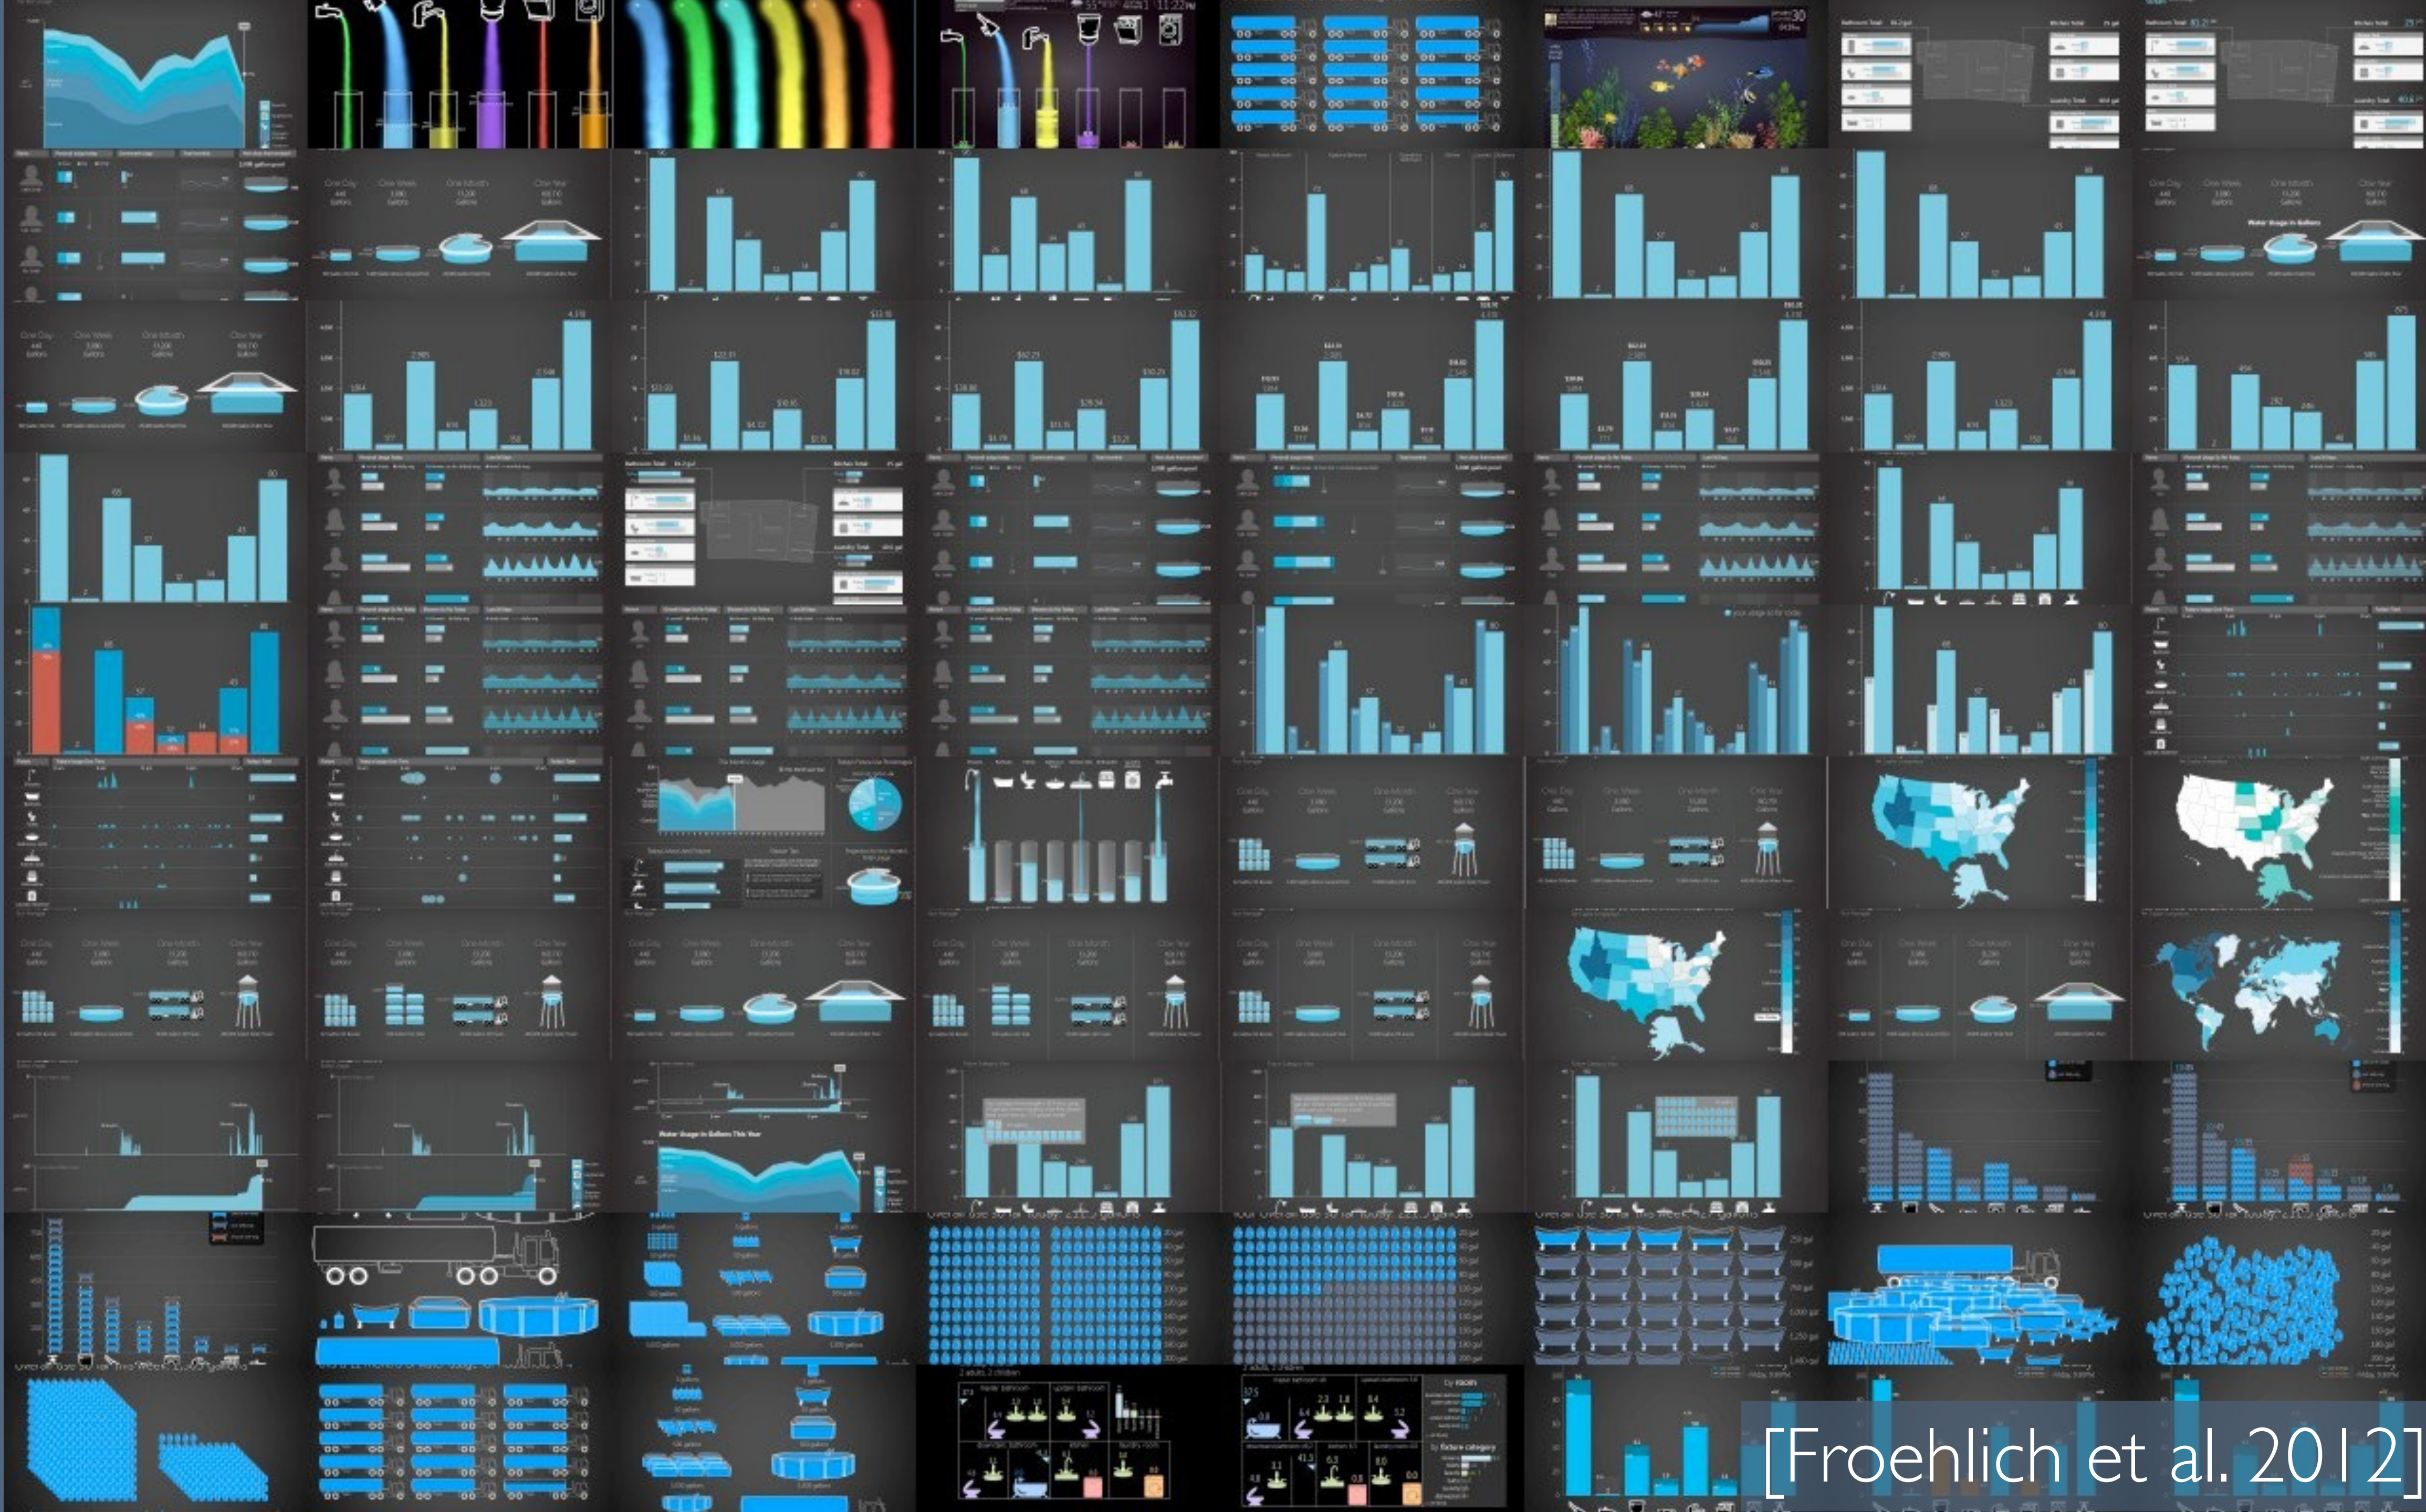


Daily average consumption by fixture for the month of May



Weekly consumption pattern for the month of May







[Froehlich et al. 2012]

Infrastructure-mediated sensing for input as well

Your house is already blanketed in wifi: we can detect minute Doppler shifts and multi-path distortions in wifi reflectance as you move
[Pu et al. 2013]

As you walk through doorways in your house, you cause momentary pressure changes in your HVAC system, enabling a recognition of where you likely are [Patel, Reynolds, and About 2008]




**Privacy and contextual
integrity**

Let's unpack this response to a ubicomp deployment

MIT Technology Review

SUBSCRIBE



SMART CITIES

Computer scientists designing the future can't agree on what privacy means

Researchers at Carnegie Mellon University wanted to create a privacy-preserving smart sensor. Then, they were accused of violating their colleagues' privacy.

By Eileen Guo & Tate Ryan-Mosley



<https://www.technologyreview.com/2023/04/03/1070665/cmu-university-privacy-battle-smart-building-sensors-mites/>

“The overall goal of this project,” Agarwal explained at an April 2021 town hall meeting, is to “build a safe, secure, and easy-to-use IoT [Internet of Things] infrastructure,” referring to a network of sensor-equipped physical objects like smart light bulbs, thermostats, and TVs that can connect to the internet and share information wirelessly.

The researchers also believe that in the long term, Mites—and building sensors more generally—are key to environmental sustainability. They see other, more ambitious use cases too. A university write-up describes this scenario: In 2050, a woman starts experiencing memory loss. Her doctor suggests installing Mites around her home to “connect to ... smart speakers and tell her when her laundry is done and when she’s left the oven on” or to evaluate her sleep by noting the sound of sheets ruffling or nighttime trips to the bathroom. “They are helpful to Emily, but even more helpful to her doctor,” the article claims.

But the Mites weren’t actually recording any video. And any audio captured by the microphones was scrambled so that it could not be reconstructed.

For some who were unhappy, exactly what data the sensors were *currently* capturing was beside the point. It didn’t matter that the project was not yet fully operational. Instead, the concern was that sensors more powerful than anything previously available had been installed in offices without consent. Sure, the Mites were not collecting data at that moment. But at some date still unspecified by the researchers, they could be. And those affected might not get a say.

Not everyone was pleased to find the building full of Mites. Some in the department felt that the project violated their privacy rather than protected it. In particular, students and faculty whose research focused more on the social impacts of technology felt that the device’s microphone, infrared sensor, thermometer, and six other sensors, which together could at least sense when a space was occupied, would subject them to experimental surveillance without their consent.


Even these sorts of adjustments wouldn’t fundamentally change how Widder feels, however. “I’m not willing to accept the premise of ... a future where there are all of these kinds of sensors everywhere,” he says.

Are these tools ethical to deploy?

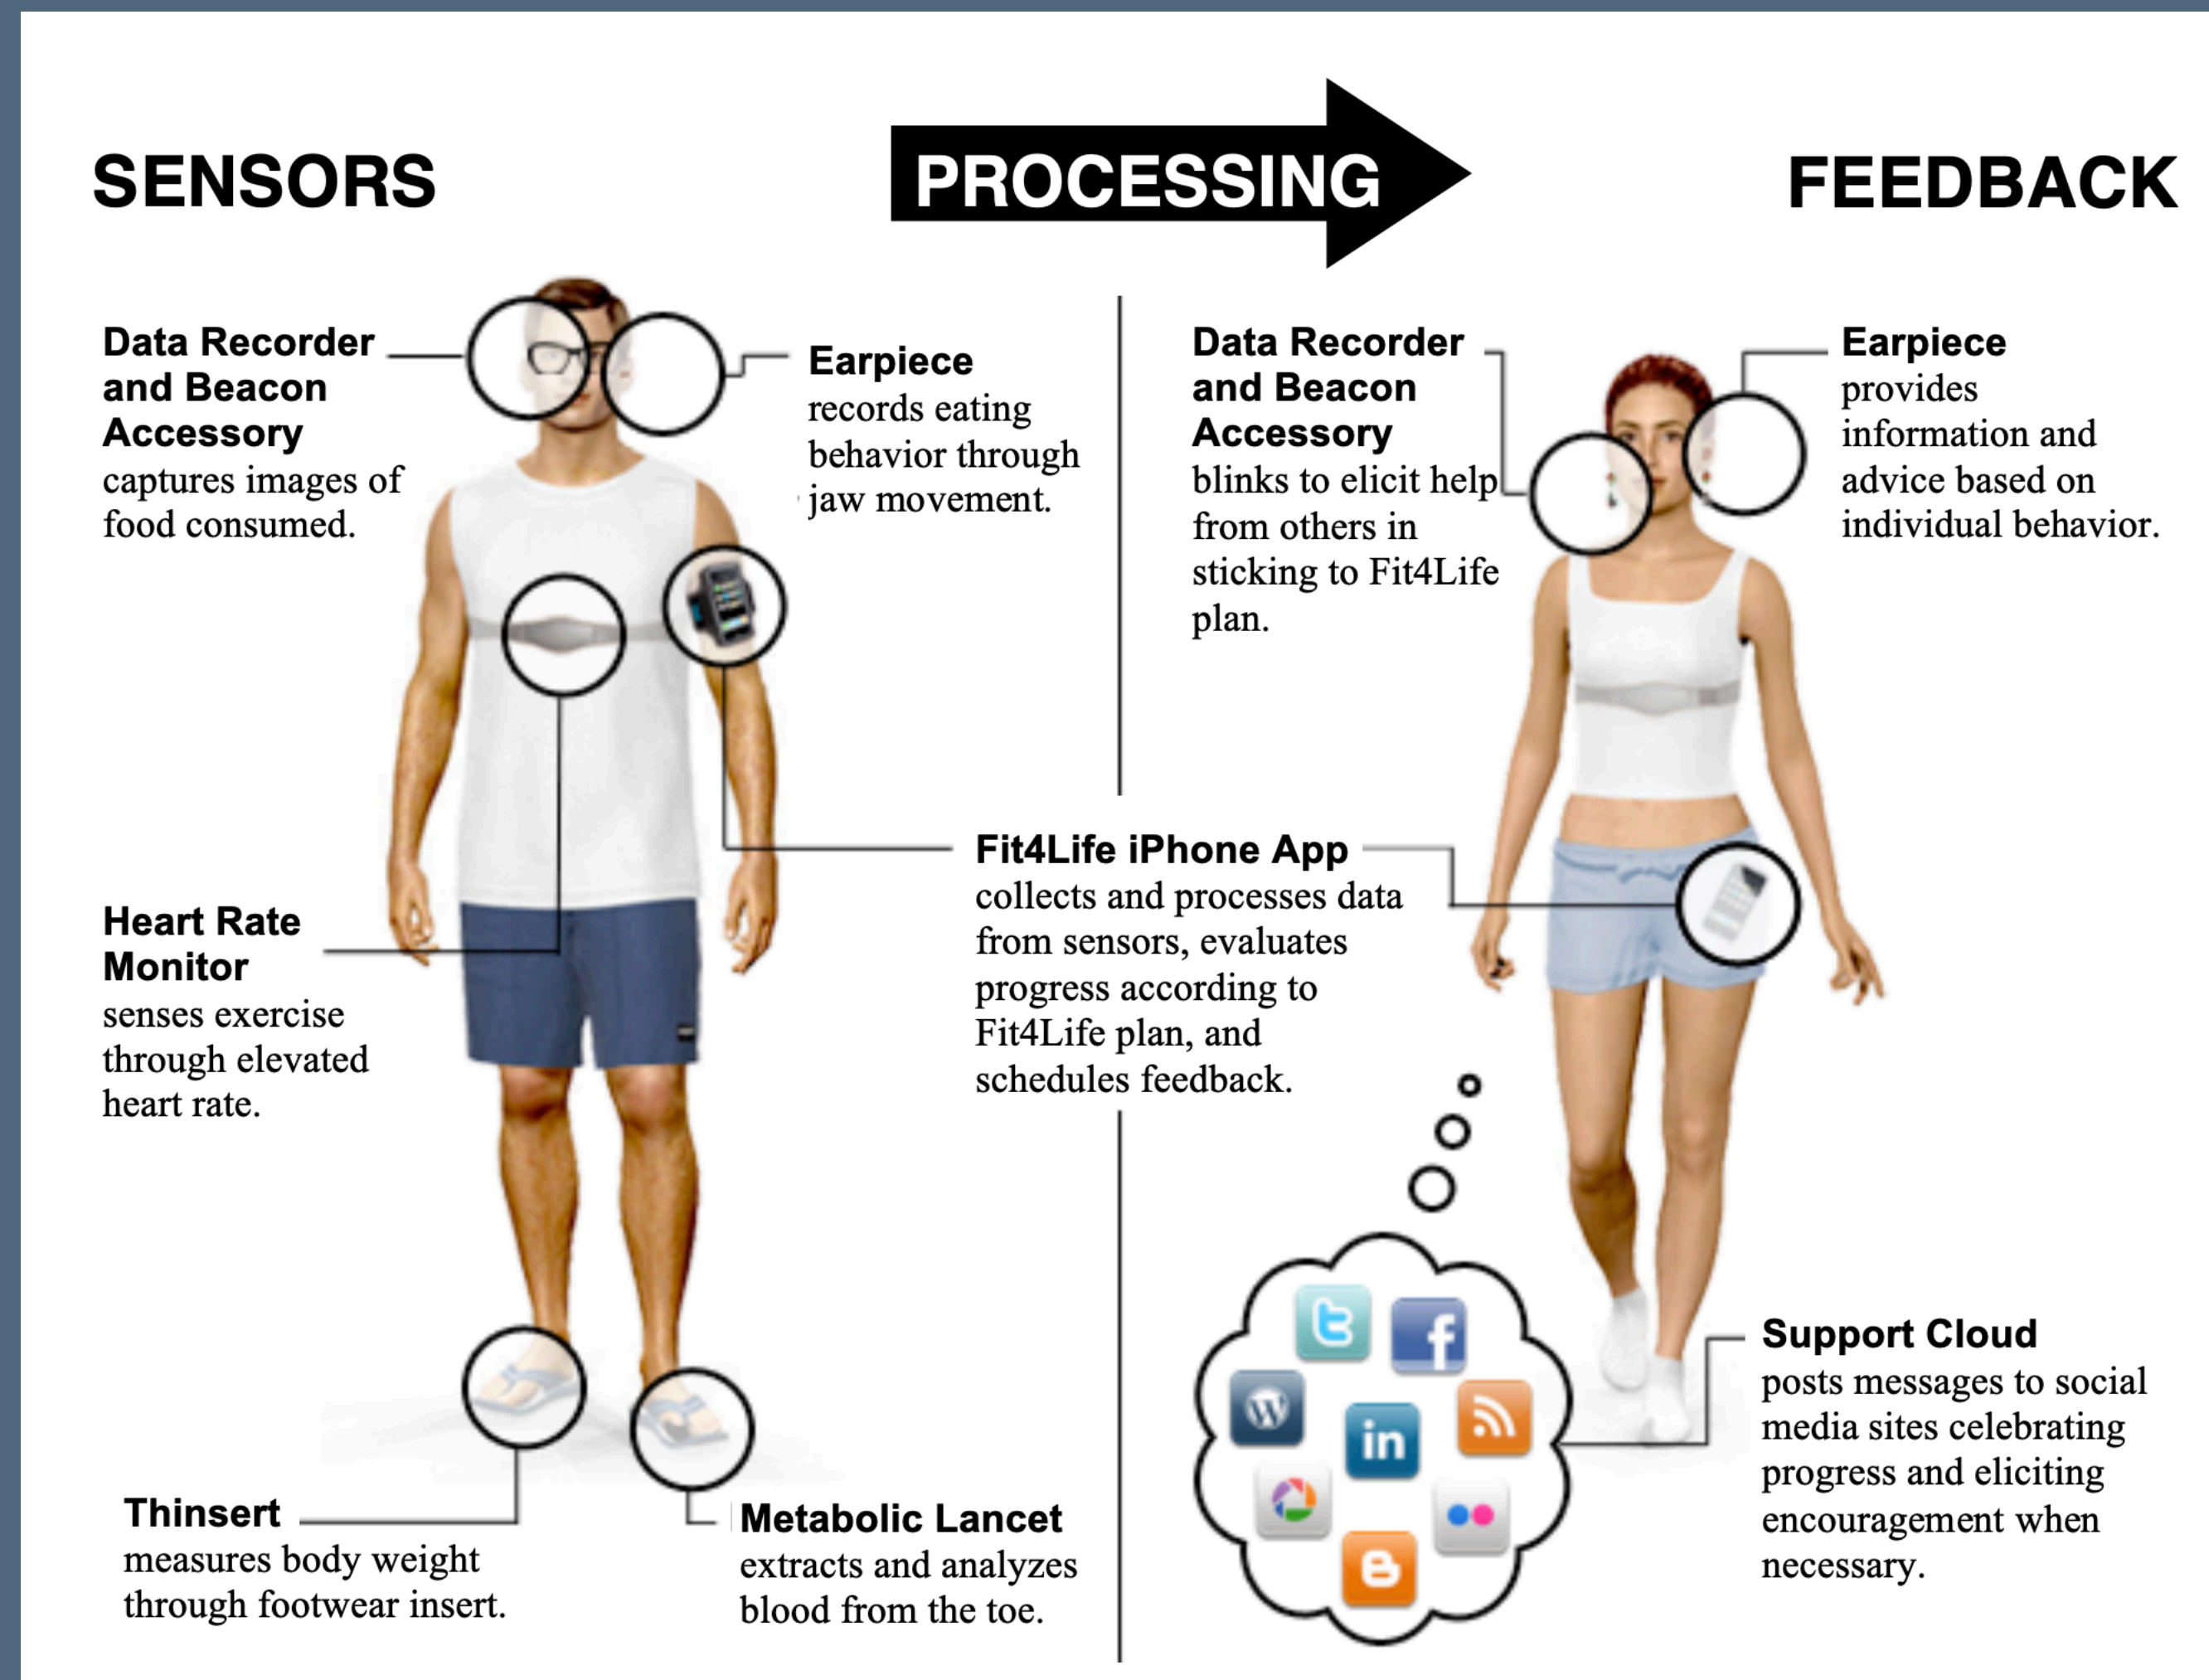
By whom?

With or without consent?

Tradeoff: privacy and autonomy,
vs. lack of access to mental
health services



Design fiction [Purpura 2011]



Are these feedback strategies **persuasion or coercion?**

Whose idea of fitness is being enacted?

Physical health coaching

GPTCoach [Jörke et al. 2024]: Use LLM to providing coaching feedback based on sensed health data

Nice to meet you, Alex!

I'm here to assist you in creating and maintaining a physical activity plan that suits your lifestyle and interests. In this program, you will be the one to design your own physical activity plan while I provide you with evidence-based support and feedback. The goal is to create a plan that reflects your preferences, resources, and goals.

This way, we can ensure that your plan is personalized, enjoyable, and sustainable for you. Do you have any questions or concerns about designing your physical activity plan or about the program in general?

I just became a new parent and am struggling to find time to exercise. I'm just not sure where to start

Alex, becoming a new parent is a major life stage and it's completely natural to feel overwhelmed. I can see why finding time for exercise could be a challenge. Remember, the key is to design a plan that fits into your current lifestyle and available resources. Don't worry, we are in this together and will find a way that works best for you. It's okay to start small, consistency is what matters the most.

Step Count

Day Week Month

< March 2024 >

Date	Step Count
Mar 3rd	22,000
Mar 4th	2,000
Mar 5th	4,500
Mar 6th	4,000
Mar 7th	8,500
Mar 8th	7,500
Mar 9th	7,000
Mar 10th	8,000
Mar 11th	8,500
Mar 12th	2,000
Mar 13th	10,000
Mar 14th	6,500
Mar 15th	8,500
Mar 16th	1,000
Mar 17th	2,000
Mar 18th	4,000
Mar 19th	3,500
Mar 20th	3,500
Mar 21st	9,000
Mar 22nd	10,500
Mar 23rd	2,000
Mar 24th	5,500
Mar 25th	1,000
Mar 26th	6,500
Mar 27th	9,500
Mar 28th	8,500
Mar 29th	3,500
Mar 30th	7,500
Mar 31st	8,000

Thank you for sharing this information, Alex. Throughout March, you averaged around 6,738 steps per day, which is a good base to start from. This suggests that even amidst your busy schedule as a new parent, you've been engaging in some activities that involve movement.

Does GPTCoach have enough context to make proper suggestions?

Privacy

Ubiquitous computing naturally raises many questions of how much privacy we are giving up in exchange for its benefits

Behavioral work has documented an empirical **privacy paradox** in which people profess to care strongly about privacy but then willingly give it up in their technology use in practice [Acquisti 2015]

Providing transparency and control are simply not enough

Contextual integrity [Nissenbaum 2004]

We often discuss information as being **private or public**.

My health data is private

My job is public (it's on my website)

My Stanford ID card RFID usage (e.g., unlocking doors) is private

But, private vs. public is **not a useful distinction** for ubicomp

Is it even possible for your Stanford RFID card usage to be fully private?

What if my smart watch sells aggregated insights to advertisers?

Is public online art usable for training generative AI?

Contextual integrity [Nissenbaum 2004]

Instead, think of information as being shared within contexts that carry specific norms—

Norms of appropriateness: What is OK and not OK to reveal in a particular context

e.g., to a doctor, it's OK to reveal medical history, less so to a bank

Norms of distribution: what is OK and not OK to share beyond the original disclosure

e.g., secrets shared with a friend are not ok to pass further; companies use the data you explicitly share with them when you sign up

Contextual integrity [Nissenbaum 2004]

Claim: there is a **privacy violation if and only if there is a violation of the norms of appropriateness or norms of flow**

Consequence: “personal information revealed in a particular context is always tagged with that context and never ‘up for grabs’”

You never get universal consent.

The Verge Menu +

DECODER

Microsoft AI chief Mustafa Suleyman says

What I was describing in that setting was the way that the world had perceived things up to that point. My take is that just as anyone can read the news and content on the web to increase their knowledge under fair use, so can an AI, because an AI is basically a tool that will help humans to learn from publicly available material. All the material that has been used for generating or training our models has been scraped from publicly available material. Where we —

VS.

engadget

AI

More than 10,500 artists sign open letter protesting unlicensed AI training

Signees include Kevin Bacon, Julianne Moore and Thom Yorke.

How would contextual integrity explain a source of friction here?

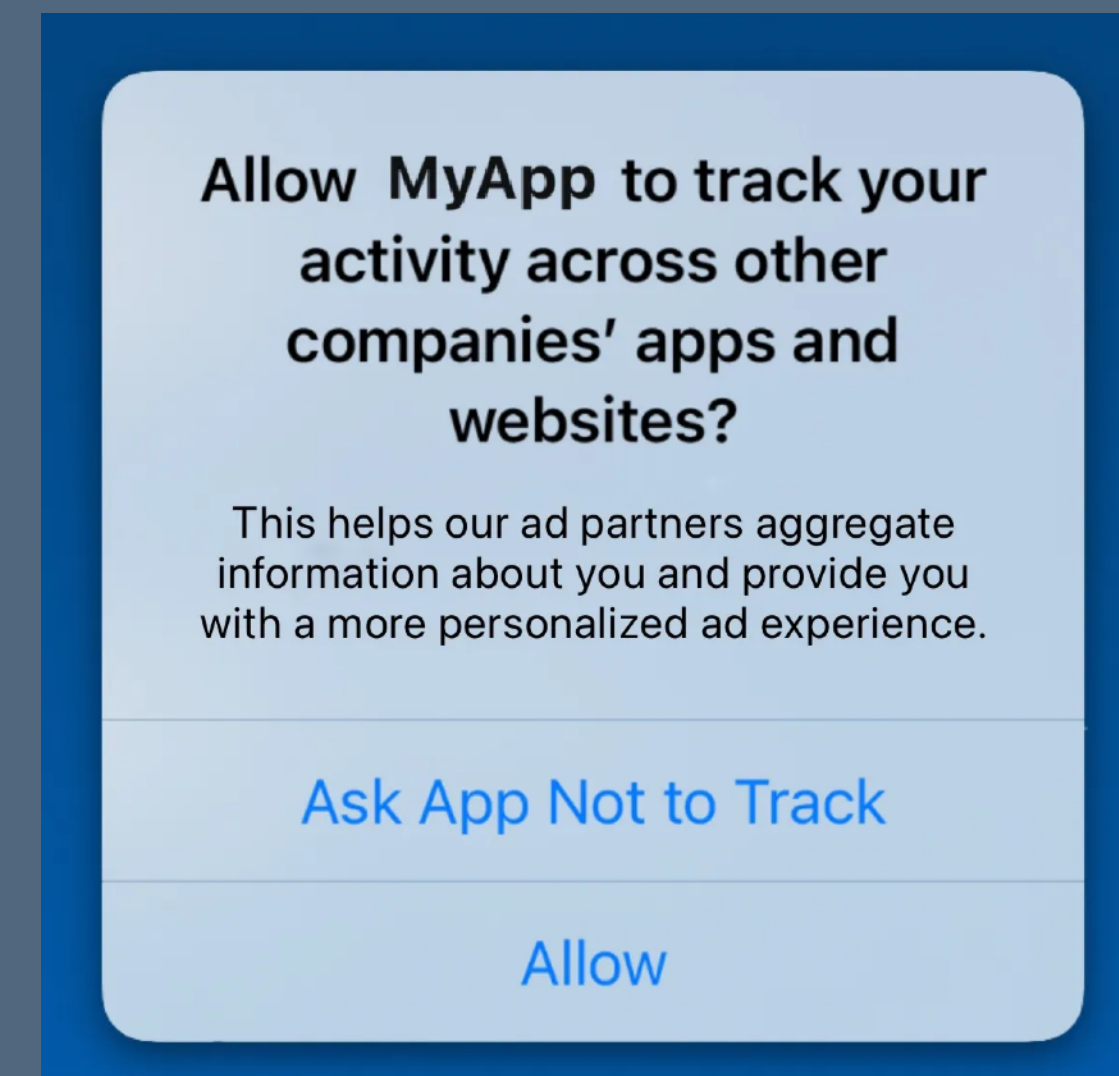
Translating contextual integrity into design

Contextual integrity diagnoses the issue, but doesn't offer many solutions that are readily amenable to design or engineering decisions

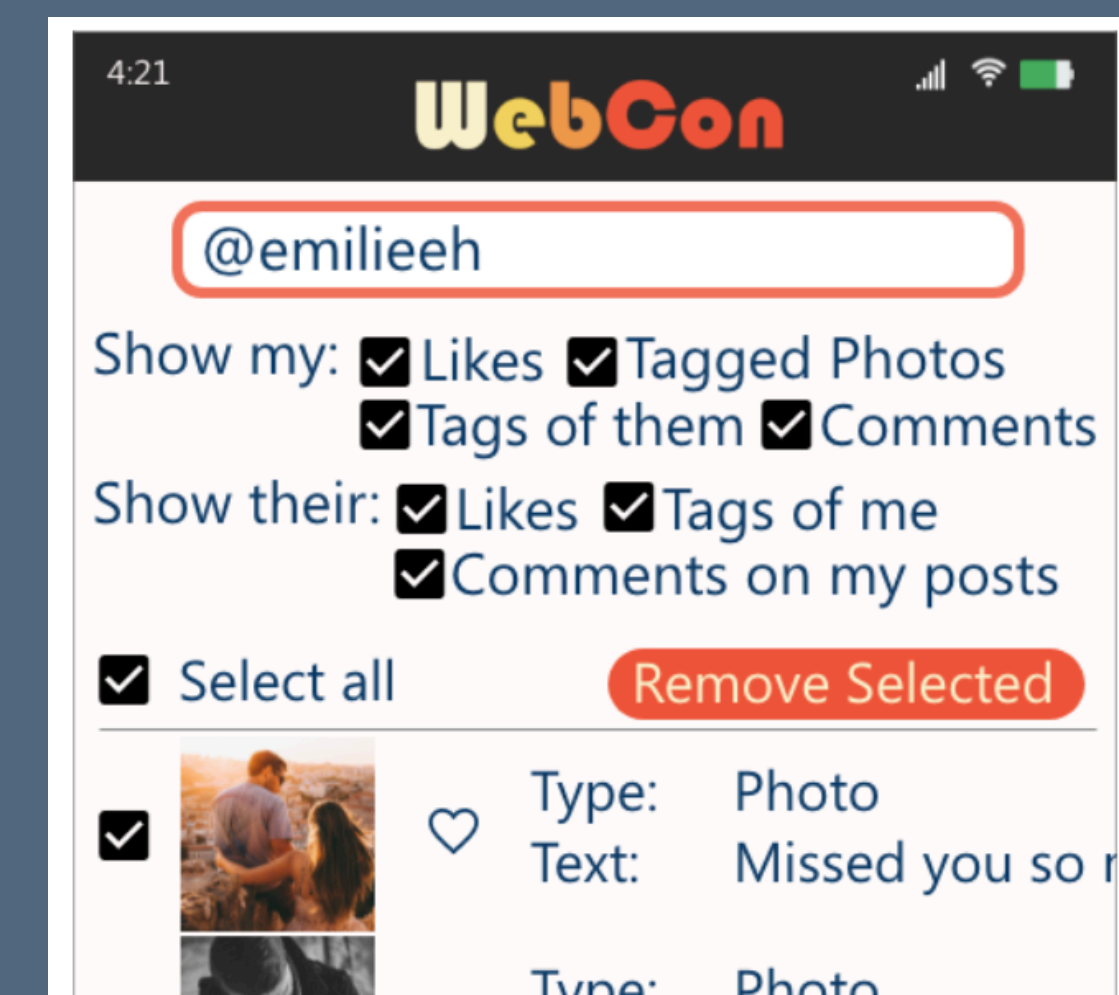
One design pattern is to require explicit opt-in 

Is it possible that LLMs could become better “contextual integrity machines?” They understand more of our norms, but will still make mistakes.

[9to5Mac]



[Im et al. 2021]



Summary

Ubicomp seeks to embed itself in **long-lived activities and goals**.

It does this across a number of domains, including: physical health, mental health and wellbeing, aging, and designing for neurodivergent populations

To achieve these goals, it seeks **noninvasive sensing** approaches

Commodity sensing: hardware each person already has or could have

Infrastructure-mediated sensing: single-point sensors that connect to existing infrastructure

Contextual integrity adds insight to the privacy questions by asking what are the norms of sharing within the sphere where the information was shared

Do's: writing a strong commentary

Do: engage with the core contributions —

Step 1 (Reflection): State the main point but then **reflect on why the ideas in the reading made sense from the authors' perspectives.**

Step 2 (Synthesis): How effectively does it convince you of that argument? **How could the argument be even more persuasive, on its own terms?**

Step 3 (Future work): What are the implications of the argument? **Given the ideas presented in the paper, what would you want to work on, or how would you modify those ideas?**

Don'ts: writing a strong commentary

Do not: stop after just disagreeing with what the authors wrote

Step 1 (Reflection): We get lots of commentaries that are mostly summary. Don't stop here. We've all read the paper. 1–2 sentences max.

Step 2 (Synthesis): It's easy to just lob criticisms and negativity. Too many commentaries are just lists of complaints. Instead, focus on: what's at the core of this idea, and why is it holding sway? What might be a better version of this idea, if you're unconvinced?

Step 3 (Future work): Too few commentaries cover this! Instead, ask yourself: what are applications of these ideas, and what follow-up ideas might be worth exploring?

References

Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein. "Privacy and human behavior in the age of information." *Science* 347.6221 (2015): 509-514.

Al Hossain, Forsad, et al. "FluSense: a contactless syndromic surveillance platform for influenza-like illness in hospital waiting areas." *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4.1 (2020): 1-28.

Ben-Zeev, Dror, et al. "CrossCheck: Integrating self-report, behavioral sensing, and smartphone use to identify digital indicators of psychotic relapse." *Psychiatric rehabilitation journal* 40.3 (2017): 266.

Bentley, Frank, et al. "Health Mashups: Presenting statistical patterns between wellbeing data and context in natural language to promote behavior change." *ACM Transactions on Computer-Human Interaction (TOCHI)* 20.5 (2013): 1-27.

Consolvo, Sunny, et al. "Activity sensing in the wild: a field trial of ubifit garden." *Proceedings of the SIGCHI conference on human factors in computing systems*. 2008.

Dey, Anind K. "Context-Aware Computing." *Ubiquitous computing fundamentals*. Chapman and Hall/CRC, 2018. 335-366.

Doryab, Afsaneh, et al. "Identifying behavioral phenotypes of loneliness and social isolation with passive sensing: statistical analysis, data mining and machine learning of smartphone and fitbit data." *JMIR mHealth and uHealth* 7.7 (2019): e13209.

Dourish, Paul. "User experience as legitimacy trap." *Interactions* 26.6 (2019): 46-49.

Dourish, Paul. "What we talk about when we talk about context." *Personal and ubiquitous computing* 8.1 (2004): 19-30

References

Escobedo, Lizbeth, et al. "MOSOCO: a mobile assistive tool to support children with autism practicing social skills in real-life situations." Proceedings of the SIGCHI conference on human factors in computing systems. 2012.

Froehlich, Jon E., et al. "HydroSense: infrastructure-mediated single-point sensing of whole-home water activity." Proceedings of the 11th international conference on Ubiquitous computing. 2009.

Froehlich, Jon, et al. "The design and evaluation of prototype eco-feedback displays for fixture-level water usage data." Proceedings of the SIGCHI conference on human factors in computing systems. 2012.

Hayes, Gillian R., et al. "Interactive visual supports for children with autism." Personal and ubiquitous computing 14.7 (2010): 663-680.

Hernandez, Javier, et al. "Under pressure: sensing stress of computer users." Proceedings of the SIGCHI conference on Human factors in computing systems. 2014.

Im, Jane, et al. "Yes: Affirmative consent as a theoretical framework for understanding and imagining social platforms." Proceedings of the 2021 CHI conference on human factors in computing systems. 2021.

Kidd, Cory D., et al. "The aware home: A living laboratory for ubiquitous computing research." International workshop on cooperative buildings. Springer, Berlin, Heidelberg, 1999.

Kientz, Julie A., et al. "Pervasive computing and autism: Assisting caregivers of children with special needs." IEEE Pervasive Computing 6.1 (2007): 28-35.

References

- Lu, Hong, et al. "Stresssense: Detecting stress in unconstrained acoustic environments using smartphones." Proceedings of the 2012 ACM conference on ubiquitous computing. 2012
- Martinez-Martin, Nicole, et al. "Ethics of digital mental health during COVID-19: crisis and opportunities." JMIR Mental Health 7.12 (2020): e23776.
- Mohr, David C., et al. "Banbury forum consensus statement on the path forward for digital mental health treatment." Psychiatric Services 72.6 (2021): 677-683
- Nandakumar, Rajalakshmi, Shyamnath Gollakota, and Jacob E. Sunshine. "Opioid overdose detection using smartphones." Science translational medicine 11.474 (2019): eaau8914.
- Nissenbaum, Helen. "Privacy as contextual integrity." Wash. L. Rev. 79 (2004): 119.
- Palen, Leysia, and Paul Dourish. "Unpacking" privacy" for a networked world." Proceedings of the SIGCHI conference on Human factors in computing systems. 2003.
- Palipana, Sameera, et al. "FallDeFi: Ubiquitous fall detection using commodity Wi-Fi devices." Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 1.4 (2018): 1-25.
- Paredes, Pablo E., et al. "Fast & furious: detecting stress with a car steering wheel." Proceedings of the 2018 CHI conference on human factors in computing systems. 2018.

References

- Patel, Shwetak N., et al. "At the flick of a switch: Detecting and classifying unique electrical events on the residential power line." International Conference on Ubiquitous Computing. Springer, Berlin, Heidelberg, 2007
- Patel, Shwetak N., Matthew S. Reynolds, and Gregory D. Abowd. "Detecting human movement by differential air pressure sensing in HVAC system ductwork: An exploration in infrastructure mediated sensing." International Conference on Pervasive Computing. Springer, Berlin, Heidelberg, 2008.
- Purpura, Stephen, et al. "Fit4life: the design of a persuasive technology promoting healthy behavior and ideal weight." Proceedings of the SIGCHI conference on human factors in computing systems. 2011.
- Thomaz, Edison, Irfan Essa, and Gregory D. Abowd. "A practical approach for recognizing eating moments with wrist-mounted inertial sensing." Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing. 2015.
- Wang, Edward Jay, et al. "Seismo: Blood pressure monitoring using built-in smartphone accelerometer and camera." Proceedings of the 2018 CHI conference on human factors in computing Systems. 2018.
- Wang, Rui, et al. "CrossCheck: toward passive sensing and detection of mental health changes in people with schizophrenia." Proceedings of the 2016 ACM international joint conference on pervasive and ubiquitous computing. 2016.
- Wang, Rui, et al. "Tracking depression dynamics in college students using mobile phone and wearable sensing." Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 2.1 (2018): 1-26.
- Xu, Xuhai, et al. "Leveraging routine behavior and contextually-filtered features for depression detection among college students." Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 3.3 (2019): 1-33.