Symposium Essays from The State of Cyberlaw: Security and Privacy in the Digital Age: AVERTING ROBOT EYES

2017

Author: Margot E. Kaminski, Matthew Rueben, William D. Smart, Cindy M. Grimm*

© 2017 Margot E. Kaminski, Matthew Rueben, William D. Smart, Cindy M. Grimm.

* Margot E. Kaminski, Associate Professor of **_Law_**, University of Colorado; Director, Silicon Flatirons Center for **_Law_**, Technology, and Entrepreneurship. Matthew Reuben, Ph.D. Candidate in Personal Robotics. William D. Smart, Associate Professor of Mechanical Engineering. Cindy M. Grimm, Associate Professor of Mechanical Engineering, Robotics Program, College of Engineering, Oregon State University.

Highlight

Abstract

Home **_robots_** will cause privacy harms. At the same time, they can provide beneficial services - as long as consumers trust them. This Essay evaluates potential technological solutions that could help home **_robots_** keep their promises, avert their eyes, and otherwise mitigate privacy harms. Our goals are to inform regulators of **_robot_**-related privacy harms and the available technological tools for mitigating them, and to spur technologists to employ existing tools and develop new ones by articulating principles for avoiding privacy harms.

We posit that home **_robots_** will raise privacy problems of three basic types: (1) data privacy problems; (2) boundary management problems; and (3) social/relational problems. Technological design can ward off, if not fully prevent, a number of these harms. We propose five principles for home **_robots_** and privacy design: data minimization, purpose specifications, use limitations, honest anthropomorphism, and dynamic feedback and participation. We review current research into privacy-sensitive robotics, evaluating what technological solutions are feasible and where the harder problems lie. We close by contemplating legal frameworks that might encourage the implementation of such design, while also recognizing the potential costs of regulation at these early stages of the technology.

Text

 [*983]

Introduction

The home is a quintessentially private location under U.S. **_law_**. The *Fourth Amendment of the U.S. Constitution* explicitly protects "houses." [1] **[*984]** The Supreme Court noted that "the very core" of the Fourth Amendment is "the right … to retreat into … [the] home" and has repeatedly protected privacy there. [2] Yet a host of technologies and new uses of technology - the Internet of Things, [3] smart grids, [4] home **_robots_** [5] - threaten that tradition of home privacy.

This Essay focuses on home **_robots_** and the privacy harms they pose. Rather than looking solely to legal solutions, it joins a growing chorus of voices emphasizing the importance of technological design. [6] Design can mitigate or prevent privacy harms. Privacy protection can be baked into a technology. We posit that if home **_robots_** are to be widely trusted, accepted, and adopted, roboticists will need to build them with privacy in mind. We aim both to inform regulators of the range of **_robot_**-related privacy harms and available technological tools for mitigating them, and to spur technologists to develop new tools by articulating principles for avoiding privacy harms.

First, we identify the types of privacy harms home **_robots_** raise. We then survey technologies that could mitigate privacy harms by home **_robots_** and outline research into privacy-sensitive robotics. We explain what types of technologies are more feasible for application to privacy-sensitive robotics in the near future, and what types are more remote. We close by returning to the **_law_**, to ask whether legal frameworks can nudge or encourage these technological adoptions - and to address the risks of trying to regulate at this relatively early stage.

If **_robots_** are to be accepted into peoples' homes, they will need to learn to give notice of surveillance, to make and keep their promises, and to avert **[*985]** their eyes. This Essay takes the first step of spelling out what this might mean with respect to **_robot_** design.

I. **_Robot_** Privacy Harms

 Home **_robots_** are coming; some are already here. They bring with them a range of privacy concerns. **_Robots_** think and act only after sensing their environment. That environmental sensing - and the storage and sharing of that information with operators and other watchers - inevitably implicates privacy. Here we describe various types of home **_robots_**, identify some prominent possible privacy harms, and briefly summarize why current **_law_** does not protect people from these harms.

A. Home **_Robots_**

---

[1] **_U.S. Const. amend. IV_** ("The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated … .").

[2] *Silverman v. United States, 365 U.S. 505, 511 (1961)* (citing Entick v. Carrington, 19 Howell's State Trials 1029, 1066 (1765); *Boyd v. United States, 116 U.S. 616, 626-30 (1886));* see also *Florida v. Jardines, 133 S. Ct. 1409, 1419 (2013); Kyllo v. United States, 533 U.S. 27, 40 (2001); Desnick v. Am. Broad. Cos., Inc., 44 F.3d 1345, 1352-53 (7th Cir. 1995)* (distinguishing *Dietemann v. Time, Inc., 449 F.2d 245 (9th Cir. 1971));* Samuel D. Warren & Louis D. Brandeis, The Right to Privacy, 4 Harv. L. Rev. 193, 202 n.1 (1890) (noting that English courts held sacred the right to privacy within the home).

[3] Andrew Guthrie Ferguson, The Internet of Things and the Fourth Amendment of Effects, *104 Calif. L. Rev. 805, 836-40 (2016);* Scott R. Peppet, Regulating the Internet of Things: First Steps Toward Managing Discrimination, Privacy, Security, and Consent, *93 Tex. L. Rev. 85, 108-12, 132-33 (2014).*

[4] See generally Andrew Paverd et al., Security and Privacy in Smart Grid Demand Response Systems, in Smart Grid Security 1 (Jorge Cuellared ed., 2014) (ebook); Andrew Paverd et al., Privacy-Enhanced Bi-Directional Communication in the Smart Grid Using Trusted Computing, 2014 IEEE Int'l Conf. on Smart Grid Comm. 872.

[5] See generally Woodrow Hartzog, Unfair and Deceptive **_Robots_**, *74 Md. L. Rev. 785 (2015);* Margot E. Kaminski, **_Robots_** in the Home: What Will We Have Agreed To?, *51 Idaho L. Rev. 661 (2015).*

[6] See Deirdre K. Mulligan & Jennifer King, Bridging the Gap Between Privacy and Design, *14 U. Pa. J. Const. L. 989 (2012);* Ira S. Rubinstein, Regulating Privacy by Design, *26 Berkeley Tech. L.J. 1409 (2011).* See also Woodrow Hartzog, Privacy's Blueprint: The Battle to Control the Design of New Technologies (forthcoming 2018).

Generally, a **_robot_** is defined by three features: it senses, thinks, and acts. [7] **_Robots_** can be independent actors or remotely operated. Both types of **_robots_** can raise privacy concerns. We here offer concrete examples of current home **_robots_**, noting that future **_robots_** are more likely to be adopted and accepted if companies can better mitigate privacy concerns.

In 2007, Bill Gates of Microsoft called for a "**_robot_** in every home." [8] Many homes, it turns out, already have one. Cleaning **_robots_** are widely available for sale, ranging from the iRobot Roomba 980, [9] to the Dyson 360 Eye **_Robot_** Vacuum, [10] the Toshiba Smarbo, [11] the Looj gutter cleaner, [12] and the Neato Botvac D3. [13] There is a cat litter box **_robot_**, the Litter-**_Robot_** III, which self-cleans and has a night light for elderly kitties. [14] Even your increasingly **[*986]** intelligent dishwasher or refrigerator could at this point be characterized as a specific-purpose **_robot_** [15]: it senses, and it acts, although it does not move. [16]

To understand why even cleaning **_robots_** raise privacy problems, it is helpful to understand how they work. There are, of course, **_robots_** explicitly designed for surveillance, like the Riley, the home-monitoring **_robot_** (advertised as giving you "eyes in the back of your head" [17]) but it is important to understand how **_robots_** designed for other purposes still implicate surveillance concerns.

---

[7] Ryan Calo, Robotics and the Lessons of Cyberlaw, _103 Calif. L. Rev. 513, 529 (2015)._

[8] Bill Gates, A **_Robot_** in Every Home, Sci. Am., Jan. 2007, at 58, _https://www.scientificamerican.com/article/a-**robot**-in-every-home/_.

[9] iRobot Roomba 980 Robotic Vacuum Cleaner, Amazon, _https://www.amazon.com/iRobot-Roomba-Robotic-Vacuum-Cleaner/dp/B013E9L4ZS_ (last visited May 17, 2017).

[10] Explore the Dyson 360 Eye, Dyson, _http://www.dyson.com/vacuum-cleaners/**robot**/dyson-360-eye.aspx_ (last visited May 17, 2017).

[11] Evan Ackerman, Toshiba Smarbo Vacuum Has Twice the Smarts, but Does It Matter?, IEEE Spectrum (Aug. 25, 2011, 10:55 AM), _http://spectrum.ieee.org/automaton/robotics/home-**robots**/toshiba-smarbo-vacuum-has-twice-the-smarts-but-does-it-matter_.

[12] iRobot Looj 330 Robotic Gutter Cleaner, Frontgate, _http://www.frontgate.com/irobot-looj-330-robotic-gutter-cleaner/531816_ (last visited May 17, 2017).

[13] Botvac D3 Connected, Neato, _https://www.neatorobotics.com/**robot**-vacuum/botvac-connected-series/botvac-d3-connected/_ (last visited May 17, 2017).

[14] Litter-**_Robot_** III Open Air, Litter-**_Robot_**, _https://www.litter-**robot**.com/litter-**robot**-iii-open-air.html_ (last visited May 17, 2017). There are some amazing reviews on Amazon for interested purchasers: "My cats DID love this, but now they just crap on the floor because it's always turning the wrong way, not emptying properly, and making loud cracking noises." Cecilia, Works Great, Until It Doesn't, Review of Litter-**_Robot_** III Open-Air - Automatic Self-Cleaning Litter Box, Amazon (Dec. 4, 2016), _https://www.amazon.com/gp/customer-reviews/RS863PYP2SFUU/ref_ =cm_cr_arp_d_rvw_ttl?ie=UTF8&ASIN=B01601QF2O.

[15] Nick Lavars, Samsung's New Smart Fridge Lets You Check in on Its Contents Through Internal Cameras, New Atlas (Jan. 6, 2016), _http://newatlas.com/samsung-family-hub-smart-fridge/41192/_ (discussing a fridge that takes a picture of its own insides whenever you close the door, which you can access on your phone when you are at the store to see what you need to buy).

[16] Erico Guizzo, Astro Teller, Captain of Moonshots at X, On the Future of AI, **_Robots_**, and Coffeemakers, IEEE Spectrum (Dec. 8, 2016, 5:39 PM), _http://spectrum.ieee.org/automaton/robotics/artificial-intelligence/astro-teller-captain-of-moonshots-at-x#gaTopicSeven_.

[17] Riley, A Smarter **_Robot_**, IndieGogo, _https://www.indiegogo.com/projects/riley-a-smarter-**robot**-drone-security_ (last visited May 17, 2017).

For example, the iRobot Roomba 980 is designed to clean a room. It navigates based on Visual Simultaneous Localization and Mapping ("VSLAM"), [18] as do the Toshiba Smarbo, the LG Hom-bot, the Neato, and the Dyson 360 Eye. [19] VSLAM works roughly as follows: the Roomba 980 has a camera on top of it that takes a picture of the environment. Then, using VSLAM software, the Roomba 980 searches for distinctive patterns, such as the edges of a couch. The **_robot_** continues to move and take images, remembering previous features and eventually building a picture-based map. To function, the **_robot_** needs to figure out its location within this map. Thus the Roomba 980 also collects its location information relative to its starting point - a process known as odometry. [20]

 [*987]   Image 1

The Roomba 980 capturing & mapping its environment. Photo: iRobot [21]

 Colin Angle, CEO and co-founder of iRobot, explains that with VSLAM, "we can create digital representations of what a home looks like so our **_robots_** can be smarter." [22] This map might include virtual walls that prevent the **_robot_** from running into things like pet food dishes or from going into locations you do not want cleaned or entered. It also means, however, that your Roomba detects when you have moved furniture, where your baby's crib is, and where you keep your safe. Your Roomba does not actually know what these items are - it knows they are obstacles - but a human looking at the images would. Moreover, if your Roomba knows where you do not want it to go, those virtual walls could indicate that something interesting or secret is behind them. That very indication is, arguably, private information. [23]

It is up for debate whether this more efficient form of navigation, which allows for single-pass cleaning of an area, actually offers an improvement over older cleaning methods, like random multi-pass cleaning. Pseudo-random cleaning methods, which older Roombas employed, may arguably be more thorough, if less efficient. [24] But the trend in cleaning **_robots_** is towards using VSLAM.

 [*988]   Robotic toys are also increasing in popularity and intelligence, and raise privacy problems. The Pleo, an adorable baby dinosaur that owners raise and train, is "able to hear, to see, to sense touch, and to detect objects" - in addition to coo endearingly. [25] Hello Barbie, released in 2015, uses a microphone to record conversations that are then transmitted to company servers, converted to text files, analyzed, and responded to, using scripted lines.

---

[18] Evan Ackerman & Erico Guizzo, iRobot Brings Visual Mapping and Navigation to the Roomba 980, IEEE Spectrum (Sept. 16, 2015, 8:30 PM), _http://spectrum.ieee.org/automaton/robotics/home-**robots**/irobot-brings-visual-mapping-and-navigation-to-the-roomba-980_.

[19] Id.

[20] Id.

[21] Id.

[22] Id.

[23] Matthew Rueben & William D. Smart, Privacy in Human-**_Robot_** Interaction: Survey and Future Work 17, 32 (2016) (unpublished   manuscript),   _http://**robots**.**law**.miami.edu/2016/wp-content/uploads/2015/07/Rueben_Smart_PrivacyInHRI_WeRobot2016.pdf_.

[24] Ackerman, supra note 11. Ackerman explains:

[A] **_robot_** vacuum can operate in one of two ways: pseudo-randomly, like a Roomba, or using a mapping pattern, like a Neato. iRobot's method involves multiple cleaning passes to clean better (maybe) at the expense of efficiency, while Neato's method covers most areas of your floor approximately once. Obviously, the Neato is much faster, so if speed is what you want, go with a vacuum that makes a map.

 Id.

[25] PLEO RB Autonomous **_Robot_** Life Form, RobotShop, _http://www.robotshop.com/en/pleo-rb-autonomous-**robot**-life-form.html_ (last visited May 17, 2017).

[26] Hello Barbie talks with and listens to your children. Robotic toys present social engineering concerns, because their owners can form attachments. Sony first sold its Aibo, a robotic dog, in 1999. When Sony retired the Aibo after its fifth generation and stopped offering repairs, owners mourned as their robotic dogs "died." [27] A human with a high degree of attachment to a robotic toy could be persuaded to buy upgrades or disclose sensitive information.

Image 2

The Pleo [28] Hello Barbie [29] The Sony Aibo [30]

Robotic toys illustrate a host of what are arguably privacy issues. They sense and record their environment and owners; they share that information **[*989]** with companies and potentially with third parties; they raise cybersecurity issues (they can be hacked) [31]; and they elicit complex emotional responses from their owners/subjects, including both affection and potentially misplaced trust. [32] Some toy _**robots**_ are explicitly advertised for their surveillance capabilities: Toys"R"Us advertises the Mebo as being able to "even spy on your family if you want him to!" [33]

Toys can be considered a subset of the broader category of social _**robots**_, which raise many of the same concerns. The Wakamaru _**robot**_ was designed by Mitsubishi to be the "world's first household _**robot**_." [34] Combining image recognition technology with technology from Mitsubishi's _**robot**_ arm division, Wakamaru was intended to "help humans by interacting." [35] Designers envisioned the _**robot**_ as being capable of doing everything from waking its owner up in the morning, to describing weather conditions and dictating email, to simply providing social companionship. Due to its rather limited capabilities, however, Wakamaru was not an overwhelming success. [36]

---

[26] James Vlahos, Barbie Wants to Get to Know Your Child, N.Y. Times Mag. (Sept. 16, 2015), _www.nytimes.com/2015/09/20/magazine/barbie-wants-to-get-to-know-your-child.html_.

[27] Chris Mills, Sony's Robotic Dogs Are Dying a Slow and Heartbreaking Death, Gizmodo (June 18, 2015), _http://gizmodo.com/sonys-robotic-dogs-are-dying-a-slow-and-heartbreaking-d-1712160637_; Jonathan Soble, A Robotic Dog's Mortality, N.Y. Times (June 17, 2015), _http://www.nytimes.com/2015/06/18/technology/robotica-sony-aibo-robotic-dog-mortality.html_.

[28] Petew, Pleo, Your Very Own Dinosaur, GadgetSpeak (Nov. 24, 2008), _https://www.gadgetspeak.com/gadget/article.rhtm/755/558783/Pleo_-_your_very_own_dinosaur.html_.

[29] Hello Barbie Doll, ToysRUs, _http://www.toysrus.com/buy/fashion-dolls/hello-barbie-doll-dkf74-71369646_ (last visited May 17, 2017).

[30] Sony Aibo Images, Jocelyn Ireson-Paine, _http://www.j-paine.org/dobbs/aibo_ images.html (last visited May 17, 2017).

[31] Samuel Gibbs, Hackers Can Hijack Wi-Fi Hello Barbie to Spy on Your Children, Guardian (Nov. 26, 2015, 6:16 AM), _https://www.theguardian.com/technology/2015/nov/26/hackers-can-hijack-wi-fi-hello-barbie-to-spy-on-your-children_.

[32] Ryan Calo, _**Robots**_ and Privacy, in Robotic Ethics: The Ethical and Social Implications of Robotics 187, 188, 195, 197 (Patrick Lin et al. eds., 2012); Hartzog, supra note 5, at 801 (pointing out that _**robots**_ may be "specifically designed to extract personal information through social engineering"); Kaminski, supra note 5, at 664 (noting that "_**robots**_' social features may elicit trust where trust is not deserved").

[33] Mebo Robotic Claw Interactive _**Robot**_, ToysRUs, _http://www.toysrus.com/product/index.jsp?productId=91846186_ (last visited May 17, 2017).

[34] World's First Practical Home-use _**Robot**_, Mitsubishi Monitor, Feb. & Mar. 2006, _https://www.mitsubishi.com/mpac/e/monitor/back/0602/story.html_.

[35] Id.

[36] Paul Miller, Mitsubishi's Wakamaru Bot Isn't Ready to Integrate into Society, Engadget (July 29, 2007), _https://www.engadget.com/2007/07/29/mitsubishis-wakamaru-bot-isnt-ready-to-integrate-into-society/_.

Image 3

Wakamaru, the "world's first household **_robot_**." [37]

  [*990] Some social **_robots_** have specifically therapeutic purposes. PARO, a robotic baby harp seal, is described by its makers as a "therapeutic **_robot_**." [38] It is intended to treat patients with Alzheimer's disease and other forms of dementia - users who may lack the ability to tell it is a **_robot_**. [39] PARO has five different kinds of sensors: tactile, light, audition, temperature, and posture sensors. It can recognize light and dark, feel strokes or strikes, and recognize the direction of a voice and words. PARO learns and remembers its name. It also remembers previous actions and tries to repeat positive actions to be petted, and eschew bad actions to avoid being struck. Hasbro has created a far less complex and less expensive therapy cat, the Joy for All Companion Pet, which responds to being petted but does not appear to gather and store as much information about its environment or user. [40]

 Image 4

PARO, a robotic baby harp seal. [41]

 More recent home **_robots_** like Pepper and JIBO are similarly envisioned as social companions but are designed for general-purpose domestic interaction. [42] JIBO uses machine learning and learns by listening to humans, and it employs facial recognition. [43] Pepper, a humanoid in design, has four directional microphones on its head, a 3D camera and two HD cameras, and an "emotion engine" whereby he can "identify your emotions by your voice as **[*991]** well as the expressions on your face." [44] The more often Pepper sees a particular face, the more accurate its detection of emotions will be. [45]

 Image 5

Pepper, a social companion **_robot_**. [46]

 Mayfield Robotics offers Kuri, a home **_robot_** with "emotive … adorable" eyes that "can even recognize faces and monitor your home when you're not there." [47] Kuri has a microphone and sound-detection technology, can learn a

---

[37] World's First Practical Home-use **_Robot_**, supra note 34.

[38] PARO Therapeutic **_Robot_**, PARO, _http://www.parorobots.com/_ (last visited May 17, 2017).

[39] Brittany A. Roston, Study Finds Robotic Paro Seal Is Therapeutic for Dementia Patients, Slash Gear (Apr. 6, 2016), _https://www.slashgear.com/study-finds-robotic-paro-seal-is-therapeutic-for-dementia-patients-06435184/_.

[40] Andy Newman, Therapy Cats for Dementia Patients, Batteries Included, N.Y. Times (Dec. 15, 2016), _http://www.nytimes.com/2016/12/15/nyregion/robotic-therapy-cats-dementia.html_. The cat, which costs $ 99.99, appears to have motion sensors and touch sensors, but it lacks audio or facial recognition. Orange Tabby Cat, Hasbro, _http://joyforall.hasbro.com/en-us/companion-cats_ (last visited May 17, 2017).

[41] Roston, supra note 39.

[42] April Glaser, Jibo is Like Alexa and a Puppy Inside One Adorable **_Robot_**, Wired (June 28, 2016), _https://www.wired.com/2016/06/jibo-like-alexa-puppy-inside-one-adorable-**robot**/_.

[43] Id.

[44] Find Out More About Pepper, SoftBank Robotics, _https://www.ald.softbankrobotics.com/en/cool-**robots**/pepper/find-out-more-about-pepper_ (last visited May 17, 2017); see also Althea Chang, At CES: **_Robots_** That Can Recognize if You're Sad, CNBC (Jan. 7, 2016),_http://www.cnbc.com/2016/01/07/at-ces-**robots**-that-can-recognize-if-youre-sad.html_.

[45] Chang, supra note 44.

[46] Id.

home's floor plan and the timing of household activities, and possesses entertainment capabilities: you "can even send her into your kids' room to tell a bedtime story." [48] Kuri's designers sought to ensure the **_robot_** was "not necessarily optimally efficient or functional, but … was approachable, calming and inviting." [49] They prioritized human interaction over efficiency, and wanted users to trust the **_robot_**, not be unnerved by it.

 [*992] Image 6

Kuri, interacting with a child. [50]

 Both therapeutic and general-purpose social **_robots_** raise similar issues to robotic toys. They record their environment, store and share information, can potentially be hacked, and deliberately engage with their owners' emotions. [51] For both robotic toys and social **_robots_**, privacy concerns can potentially prevent widespread adoption.

Take Hello Barbie as an example. Privacy concerns about the doll received considerable press attention. [52] One article asked, "Is Hello Barbie eavesdropping on your kids?" [53] Another explained that hackers could hijack Barbie and use her for surveillance. [54] The Campaign for a Commercial Free Childhood launched a campaign against Hello Barbie, explaining that "kids using "Hello Barbie' won't only be talking to a doll, they'll be talking directly to a toy conglomerate whose only interest in them is financial." [55] The toy ultimately did not perform as hoped, generating publicity but not the hoped-for sales. [56] While the poor sales could also be attributed to other glitches like **[*993]** a "shaky" Internet connection, [57] privacy concerns likely played a part. Privacy concerns may also

---

[47] Life with Kuri, Kuri, *https://www.heykuri.com/living-with-a-personal-robot* (last visited May 17, 2017).

[48] Id.

[49] Darrell Etherington, Home **_Robot_** Kuri Is like an Amazon Echo Designed by Pixar, TechCrunch (Jan. 3, 2017), *https://techcrunch.com/2017/01/03/home-robot-kuri-is-like-an-amazon-echo-designed-by-pixar/*.

[50] Life with Kuri, supra note 47.

[51] Elaine Sedenberg, John Chuang & Deirdre Mulligan, Designing Commercial Therapeutic **_Robots_** for Privacy Preserving Systems and Ethical Research Practices Within the Home, 8 Int'l J. Soc. Robotics 575, 575-79 (2016).

[52] See, for example, the following commentaries: Lori Andrews, Hello Barbie, Goodbye Privacy, Chi. Trib. (Nov. 27, 2015, 4:18 PM), *http://www.chicagotribune.com/news/opinion/commentary/ct-hello-barbie-privacy-icloud-perspec-20151127-story.html*, and Sarah Halzack, Privacy Advocates Try to Keep "Creepy,' "Eavesdropping' Hello Barbie from Hitting Shelves, Wash. Post (Mar. 11, 2015), *https://www.washingtonpost.com/news/the-switch/wp/2015/03/11/privacy-advocates-try-to-keep-creepy-eavesdropping-hello-barbie-from-hitting-shelves/*.

[53] Susan Linn, Is Hello Barbie Eavesdropping on Your Kids?, CNN (Dec. 4, 2015), *http://www.cnn.com/2015/12/04/opinions/linn-hello-barbie-privacy/*.

[54] Gibbs, supra note 31.

[55] Leigh Weingus, Talking Barbie Could Eavesdrop on Kids, Critics Warn, Huff. Post (Mar. 11, 2015), *http://www.huffingtonpost.com/2015/03/11/wifi-barbie_n_6847736.html* (quoting Campaign for a Commercial Free Childhood, Stop Mattel's "Hello Barbie" Eavesdropping Doll, *http://www.commercialfreechildhood.org/action/stop-mattel%E2%80%99s-hello-barbie-eavesdropping-doll* (last visited May 17, 2017)).

[56] Matthew Townsend, Hello Barbie Pleads "Buy Me' As Mattel Doll Fails to Catch Fire, Bloomberg Tech. (Apr. 20, 2016), *https://www.bloomberg.com/news/articles/2016-04-20/hello-barbie-pleads-buy-me-as-mattel-doll-fails-to-catch-fire*.

[57] Paul R. La Monica, Nobody Puts Barbie in a Corner. Mattel Soars, CNN Money (Feb. 2, 2016), *http://money.cnn.com/2016/02/02/investing/mattel-earnings-barbie/* (describing Hello Barbie as "a little bit creepy"). For context, La Monica's article was written when Mattel announced growth around the holidays in 2015. See Matthew Townsend, Barbie's Holiday Sales Grow for First Time in Four Years, Bloomberg (updated Feb. 2, 2016, 4:20 PM), *https://www.bloomberg.com/news/articles/2016-02-01/mattel-earnings-top-estimates-as-barbie-shows-signs-of-rebound*;

preclude distribution to markets where privacy is taken more seriously. For example, Germany recently banned the "My Friend Cayla" doll, an interactive toy that allegedly transmits recorded conversations to a voice recognition company in the United States, citing privacy concerns. [58]

Joining a growing cadre of scholars that emphasizes the necessity of trust in the information age, we contend home **_robots_** will be widely adopted only if users can trust that their information will not be illicitly gathered, shared, or misused. [59] Amazon appears to agree: it recently filed a motion to quash a search warrant directed at obtaining information gathered by the Amazon Echo, arguing that "rumors of an Orwellian federal criminal investigation into the reading habits of Amazon's customers could frighten countless potential customers." [60]

B. **_Robot_** Privacy Harms

 What do we mean when we talk about **_robot_** privacy harms? Privacy has been defined in a wide range of ways, from control, to decisional autonomy, to complex taxonomies of privacy violations. [61] Home **_robots_** potentially raise a number of different privacy harms. As illustrated above, **_robots_** **[*994]** gather, process, share, and store information. While not intending to be exhaustive, we posit that **_robots_** raise privacy problems of three basic types: (1) data privacy problems; (2) boundary management problems; and (3) social/relational problems. While **_robots_** also implicate cybersecurity concerns, those are largely outside the scope of this Essay, which focuses on harms arising from normal operation without the intervention of malicious parties. [62]

First, **_robots_**, like a wide range of sensor-laden technologies, raise a host of data privacy concerns. [63] These concerns are provoked by the gathering, sharing, and storage of both sensitive information and information from which sensitive inferences can be drawn. Concerns are also raised by sharing information out of context in a way that functionally makes that information sensitive. [64]

Perhaps most obviously, **_robots_** may directly gather recognizably sensitive information from the home environment. [65] A **_robot_** might photograph a prescription bottle or overhear a conversation about a cancer diagnosis. It might record your child's conversation or capture a picture of your latest bank balance. United States federal privacy **_law_** is largely organized around protecting different classes of sensitive information - health information, financial

---

Townsend, supra note 56 ("Hello Barbie generated the most buzz the toymaker has received in years, if not decades. But for the most part, it's been a dud.").

[58] Soraya Sarhaddi Nelson, Germany Bans "My Friend Cayla' Doll Over Spying Concerns (NPR radio broadcast Feb. 20, 2017, 4:40 PM), _http://www.npr.org/templates/transcript/transcript.php?storyId=516292295_.

[59] See Jack Balkin, Information Fiduciaries and the First Amendment, _49 U.C. Davis L. Rev. 1183, 1187 (2016)_ ("My goal, in other words, is to shift the focus of the First Amendment arguments about privacy from the kind of information to the kinds of relationships - relationships of trust and confidence - that governments may regulate in the interests of privacy."); Kiel Brennan-Marquez, Fourth Amendment Fiduciaries, _84 Fordham L. Rev. 611, 613 & n.5 (2015);_ Neil Richards & Woodrow Hartzog, Taking Trust Seriously in Privacy **_Law_**, 19 Stan. Tech. L. Rev. 431, 432, 434 (2016).

[60] Memorandum of **_Law_** in Support of Amazon's Motion to Quash Search Warrant at 14, State of Arkansas v. James A. Bates, No. CR-2016-370-2 (Cir. Ct. Ark. Feb. 17, 2017 (quoting In Re Grand Jury Subpoena to Amazon.com dated **_August 6, 2007, 246 F.R.D. 570, 573 (W.D. Wis. 2007))._** This motion has since been dropped, as the suspect agreed to hand over recordings. Eliott C. McLaughlin, Suspect OKs Amazon to Hand Over Echo Recordings in Murder Case, CNN (Apr. 26, 2017, 2:52 PM), _http://www.cnn.com/2017/03/07/tech/amazon-echo-alexa-bentonville-arkansas-murder-case/_

[61] Daniel J. Solove, Understanding Privacy 10-11 (2008); Rueben & Smart, supra note 23, at 4-7.

[62] See generally Derek E. Bambauer, Privacy Versus Security, _103 J. Crim. L. & Criminology 667 (2013)._

[63] See Sedenberg, Chuang & Mulligan, supra note 51, at 578 (discussing "Aspects of Informational Privacy Specific to Therapeutic **_Robots_**").

[64] Id. at 580 (discussing contextual integrity).

[65] Paul Ohm, Sensitive Information, _88 S. Cal. L. Rev. 1125, 1128-29 (2015)._

information, information about children - and those classes of information are likely to be implicated by recordings of the home. As discussed below, however, U.S. privacy *law* is also largely limited to coverage of particular entities or particular relationships. If you share your health information with your doctor, she has privacy obligations. If you share it with Google, such legal obligations likely do not attach. [66] Users who expect legal protection for sensitive information revealed to home *robots* may be unpleasantly surprised to find themselves without legal recourse.

Another classic data privacy concern arises from storage and analysis of bulk amounts of information. Like most digital technology, home *robots* might "remember" - or really, store - information far longer than humans are calibrated to understand. One of the ways in which we manage our social relationships is by spacing them out over time or by depending on the faulty memories of the people with whom we interact. [67] By introducing temporal **[*995]** permanence into the home environment, home *robots* may collapse those interactions and eliminate important relationship-management tools.

Storing information allows for data analysis, which also allows inferences to be drawn about repeated behaviors. [68] By recording large amounts of information, companies may be equipped to know what time daily showers happen or when homeowners are typically out of the house. They might infer illness, vulnerability, or other changes in physical or emotional well-being that make users more susceptible to advertising or particular marketing appeals. [69] They might figure out and classify a homeowner's religion or race, for potential use in targeted advertising. [70] *Robot* users may fail to understand that they are in fact communicating sensitive information about themselves in these interactions. When you interact with your dog, it does not remember everything, discern all patterns, or communicate sensitive information to others. When you interact with your *robot* dog, it may do all of these things.

Robots may, like other information technology, enable individuals or companies to take information that has been shared in one context and share or use it in another. For example, if you ask your *robot* (or your search engine) about the symptoms of an illness, you do not expect that information to be shared with, say, your employer or even online advertisers. Helen Nissenbaum has articulated a theory of privacy as "contextual integrity" that looks at information flows and their disruption. [71] *Robots* may threaten privacy by threatening contextual integrity. [72] They may fail to safeguard information within a particular context and relatedly fail to alert users to the possibility that the context in which they share information may not be what it appears to be. What looks like your empty living room may, with a *robot* present, in fact be a company server, a behavioral advertiser, or even a government agency.

---

[66] To Whom Does the Privacy Rule Apply and Whom Will It Affect?, Nat'l Insts. of Health, *https://privacyruleandresearch.nih.gov/pr_06.asp* (last visited May 17, 2017).

[67] Viktor Mayer-Schonberger, Delete: The Virtue of Forgetting in the Digital Age (2011); Margot E. Kaminski, Regulating Real-World Surveillance, *90 Wash. L. Rev. 1113, 1134 (2015).*

[68] Fed. Trade Comm'n, Big Data: a Tool for Inclusion or Exclusion? 1 (2016), *https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf*.

[69] Sarah Gray, One Woman's Attempt to Hide Her Pregnancy from Big Data - It's More Difficult Than You'd Expect, Salon (Apr. 28, 2014), *http://www.salon.com/2014/04/28/one_womans_attempt_to_hide_her_pregnancy_from_big_data/*; see also Sedenberg, Chuang & Mulligan, supra note 51, at 579 (discussing how therapeutic *robots* might "infer psychological and mental states").

[70] Sapna Maheshwari & Mike Isaac, Facebook Will Stop Some Ads from Targeting Users by Race, N.Y. Times (Nov. 11, 2016), *http://www.nytimes.com/2016/11/12/business/media/facebook-will-stop-some-ads-from-targeting-users-by-race.html*.

[71] Helen Nissenbaum, Privacy in Context: Technology, Policy, and the Integrity of Social Life 127, 128 (2010). The Obama administration largely adopted this theory of privacy in a 2014 report encapsulated in the proposed, but not enacted, Consumer Privacy Bill of Rights. Exec. Office of the President, Big Data: Seizing Opportunities, Preserving Values 19 (2014) ("Respect for Context: Consumers have a right to expect that organizations will collect, use, and disclose personal data in ways that are consistent with the context in which consumers provide the data.").

[72] See Sedenberg, Chuang & Mulligan, supra note 51, at 580.

[*996] The second class of **_robot_** privacy harms involve boundary management problems. [73] **_Robots_** might see through or move around barriers humans use to manage their privacy, or they might "see" things using senses humans would not know to guard against. [74] Social psychologist Irwin Altman theorized that privacy is the dynamic process of managing one's social accessibility using the tools of one's environment. [75] As one of us has discussed elsewhere at length, this model of privacy fits well with real-world surveillance concerns. [76] Where with no **_robot_** a person might rely on the walls of the rooms of her house to safeguard her privacy in a particular room, a **_robot_** that uses, say, thermal sensors may be able to "see" through the walls and render them ineffective. [77] Even a **_robot_**'s ability to move around physical barriers - which your computer, for example, cannot do - renders those barriers less effective as a tool for privacy management.

Robots may use sensors that humans do not expect to have to guard against, or are just not equipped to avoid. [78] We may be good, or at least better, at guarding against visual surveillance, since we can generally see those who see us and adapt accordingly. [79] If we see a **_robot_** staring at us in one room, we might move to another one. But we may incorrectly assume that **_robots_**, especially those designed to be anthropomorphized, have similar sensory abilities to humans. Other sensors - from audio sensors to thermal imaging - may threaten our ability to manage our accessibility-slash-privacy because we are less well equipped, or not equipped at all, to guard against them. Looking at an adorable **_robot_**, we may forget they have radar and thermal sensors, and they are constantly sniffing the data packets coming from our phones. Similarly, we may miscalibrate how audible we are to a **_robot_** in another room, expecting that it has human hearing levels when in fact it is capable of listening in with, for example, a laser Doppler vibrometer, which can hear a heartbeat at 300 yards. [80] And we cannot make ourselves less visible to a thermal imager or know what we are revealing by moving the **_robot_** from place to place. Home **_robots_** thus threaten privacy by recording using senses humans are not prepared to address.

[*997] The third category of privacy problems **_robots_** raise are the social/relational problems raised by designing human-**_robot_** interactions against the backdrop of social behavior. As discussed above, one of the more unique aspects of **_robots_** compared to other information technologies is their potential to develop social relationships with humans - or at least, to make humans feel and behave like a relationship exists. [81] This has significant implications for privacy. If you trust a **_robot_**, you might disclose more. [82] You may feel like you are talking to your dog or friend when in fact you are talking to a corporation. [83] Research also shows that the perceived persona of a **_robot_** can

---

[73] See generally Irwin Altman, The Environment and Social Behavior (1975); Kaminski, supra note 67.

[74] Kaminski, supra note 5, at 661-63; Calo, supra note 32, at 192.

[75] Altman, supra note 73.

[76] Kaminski, supra note 67, at 1113, 1117.

[77] _Kyllo v. United States, 533 U.S. 27, 36 (2001);_ see also Kaminski, supra note 67, at 1119.

[78] Marc J. Blitz, The Right to Map (and Avoid Being Mapped): Reconceiving First Amendment Protection for Information Gathering in the Age of Google Earth, _14 Colum. Sci. & Tech L. Rev. 115, 190 (2012)._

[79] Rueben & Smart, supra note 23, at 34.

[80] Id.

[81] Calo, supra note 7, at 545; Kate Darling, "Who's Johnny?' Anthropomorphic Framing in Human-**_Robot_** Interaction, Integration, and Policy, in **_Robot_** Ethics 2.0 (Patrick Lin et al. eds., forthcoming 2017), _https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2588669_.

[82] Calo, supra note 32, at 187, 188, 195, 197; Hartzog, supra note 5, at 801; Kaminski, supra note 5, at 664.

[83] Weingus, supra note 55.

really matter; people may trust **_robots_** more if they perceive them as specialists, [84] and disclose different kinds of information based on perceived gender. [85]

Even failed design has privacy implications. **_Robots_** may fail to observe personal space. [86] Research has shown that people react to the invasion of personal space by artificial agents (that is, images of people) much as they do to invasions by actual people. [87] Similarly, **_robots_** are known for having a "constant gaze" problem: a constant stare that makes humans deeply uncomfortable. [88] These in-person perceived privacy invasions may affect peoples' behavior in the home, even if they do not reflect downstream information harms. When designers try to work around these problems, this can also enable information privacy harms. For example, designers are learning to lower **_robot_** eyes and keep **_robots_** out of your personal space, even as the machine is still actually recording. This can send inaccurate messages to **_robot_** users, who may miscalibrate their behavior accordingly.

Robots' failures to respond to social cues implicate information privacy in another way: social cues become a less effective method for managing social accessibility. While a visitor to your home may understand from both broader social context and your specific actions that it is impermissible to record your child, or enter a bedroom, or go onto your computer, a **_robot_** **[*998]** will not read its environment in the same way. You can and do control visitor behavior through social norms and social cues; your guidance to a **_robot_** will need to be more explicit and of a different kind.

Robots shield their operators from both social cues and social sanctions. Where a person may receive immediate social sanction for taking out their phone to record your child against your wishes, a **_robot_**'s operator, whether an individual or a company, is remote from that moment. This creates a problem of dissociation for **_robot_** operators: the de-coupling of one's body from one's actions, making social signals that can already be challenging to discern in in-person interactions even more difficult or impossible to read remotely. [89] Even well-meaning actors can be disinhibited by dissociation, and bad actors can no longer be directly shamed into behaving.

All of the above harms can have larger implications for society. A lack of privacy can cause conformity and chilling effects. [90] A loss of solitude in the home might have other psychological effects, removing an important reprieve from the busy world. [91] **_Robot_** privacy harms will implicate all of the values implicated by privacy harms: autonomy, dignity, fairness, trust and intimacy, trust and sociality, democratic participation, and more. [92]

C. Why Current **_Law_** Is Not Enough

---

[84] Rueben & Smart, supra note 23, at 2.

[85] Laura Dattaro, Bot Looks like a Lady, Slate (Feb. 4, 2015, 1:12 PM), http://www.slate.com/articles/technology/future_tense/2015/02/_robot_gender_is_it_bad_for_human_women.html_; Aaron Powers et al., Eliciting Information from People with a Gendered Humanoid **_Robot_**, 14 IEEE Int'l Workshop on **_Robot_** & Human Interactive Comm. 158, 158-59 (2005).

[86] Rueben & Smart, supra note 23, at 26.

[87] Id. at 2.

[88] Id. at 24-25.

[89] Id. at 16.

[90] Margot E. Kaminski & Shane Witnov, The Conforming Effect: First Amendment Implications of Surveillance, Beyond Chilling Speech, _49 U. Rich. L. Rev. 465, 467 (2015);_ Jonathon W. Penney, Chilling Effects: Online Surveillance and Wikipedia Use, _31 Berkeley Tech. L.J. 117, 117 (2016)._

[91] Calo, supra note 32, at 196 ("Privacy provides "a respite from the emotional stimulation of daily life' that the presence of others inevitably engenders." (quoting Alan F. Westin, Privacy and Freedom 35 (1967)).

[92] See, e.g., Valerie Steeves, Reclaiming the Social Value of Privacy, in Lessons from the Identity Trail: Anonymity, Privacy, and Identity in a Networked Society 191, 196-98 (Ian Kerr et al. eds., 2009) (discussing Westin's theories of personal autonomy, emotional release, self-evaluation, and limited and protected communication).

Current U.S. _law_ does not adequately protect against these privacy harms. For one, there is no general federal data protection _law_ in the United States, only industry-specific protection for particular kinds of information, usually hinging on particular relationships or targeting particular entities. For example, the Health Insurance Portability and Accountability Act ("HIPAA") protects individually identifiable health information under the Privacy Rule, but the Act applies only when that information is handled by "covered entities" or the business associates of covered entities. [93] These entities are defined as health plans, health care clearinghouses, and health care **[*999]** providers. [94] HIPAA does not cover most researchers, employers, life insurers, schools, or many others. [95]

Similarly, the Fair Credit Reporting Act ("FCRA") governs credit information, but only when it is handled by "consumer reporting agencies." [96] Courts and government agencies have interpreted the term for the digital age to include data brokers that profile consumers and sell those profiles, but those broader interpretations have been challenged. [97] Even the Children's Online Privacy Protection Act ("COPPA"), which protects personally identifiable information about children under thirteen, applies only to websites "directed to children" or that have actual knowledge that they are collecting personal information from children. [98] The COPPA Rule sets out factors for determining whether a website is directed to children, including subject matter, content, the use of child-oriented activities and incentives, advertising directed at children, and more. [99] Thus, those U.S. privacy _laws_ that protect particularly sensitive information may not be prepared for the new entities involved in handling sensitive information gathered by home _robots_. [100]

United States privacy _law_ is challenged by home _robots_ in other ways. The Federal Trade Commission ("FTC") has in many ways become the de **[*1000]** facto federal privacy regulator. [101] Under its Section 5 authority, the

---

[93] To Whom Does the Privacy Rule Apply and Whom Will It Affect?, supra note 66; see also Sedenberg, Chuang & Mulligan, supra note 51, at 578-79 (discussing HIPAA's coverage and lack thereof with respect to therapeutic _robots_).

[94] To Whom Does the Privacy Rule Apply and Whom Will It Affect?, supra note 66.

[95] Your Rights Under HIPAA, Dep't Health & Human Servs., _https://www.hhs.gov/_ hipaa/for-individuals/guidance-materials-for-consumers/ (last visited May 17, 2017).

[96] _15 U.S.C. § 1681a_(f) (2012). The statute provides:

The term "consumer reporting agency" means any person which, for monetary fees, dues, or on a cooperative nonprofit basis, regularly engages in whole or in part in the practice of assembling or evaluating consumer credit information or other information on consumers for the purpose of furnishing consumer reports to third parties, and which uses any means or facility of interstate commerce for the purpose of preparing or furnishing consumer reports.

Id.

[97] Press Release, Fed. Trade Comm'n, Spokeo to Pay $ 800,000 to Settle FTC Charges Company Allegedly Marketed Information to Employers and Recruiters in Violation of the FCRA (June 12, 2012), _https://www.ftc.gov/news-events/press-releases/2012/06/spokeo-pay-800000-settle-ftc-charges-company-allegedly-marketed_.

[98] _15 U.S.C. § 6501_(10)(A) (2012). This statute provides: "The term "website or online service directed to children' means - (i) a commercial website or online service that is targeted to children; or (ii) that portion of a commercial website or online service that is targeted to children." Id. COPPA does not apply to websites that merely link to other websites targeted to children. See _15 U.S.C. § 6501_(10)(B) (2012); see also _16 C.F.R.§§312.2_, _312.3 (2016)_; Sedenberg, Chuang & Mulligan, supra note 51, at 583 (discussing COPPA's limited application to therapeutic _robots_).

[99] _16 C.F.R. § 312.2 (2016)_; see also Complying with COPPA: Frequently Asked Questions, Fed. Trade Comm'n, _https://www.ftc.gov/tips-advice/business-center/guidance/complying-coppa-frequently-asked-questions#Web%20sites%20and%20online_ (last updated Mar. 20, 2015).

[100] See B.J. Ard, Confidentiality and the Problem of Third Parties: Protecting Reader Privacy in the Age of Intermediaries, _16 Yale J.L. & Tech. 1 (2013)_ (explaining that reader privacy _laws_ fail to protect readers in the digital age because they target institutions such as libraries, rather than protect the act of reading).

[101] Daniel J. Solove & Woodrow Hartzog, The FTC and the New Common _Law_ of Privacy, _114 Colum. L. Rev. 583, 585 (2014)._

FTC protects consumers from unfair and deceptive practices, including poor privacy and data protection policies. The FTC may well be equipped to address a variety of the above-mentioned privacy harms. [102] But the FTC customarily governs a relationship between consumers and the company from which they purchase something, or with which they have an agreement. This approach is ill-suited to protecting the privacy of third parties (non-owners) impacted by products that move in the real world, including guests in your home, children in your home, or your neighbor caught on camera by your lawn-mowing **_robot_**. As Meg Jones has described it, the FTC is not necessarily equipped to handle the privacy problems raised by the "Internet of Other People's Things." [103]

State privacy **_law_** also may not cover the problems discussed here. Many states have attorney generals that enforce privacy via consumer protection **_laws_**, [104] but that enforcement, like FTC enforcement, customarily hinges on a consumer relationship with a company, which will not extend to third parties impacted by surveillance. States have privacy torts, including intrusion upon seclusion and public disclosure of private facts, but these torts struggle to capture information privacy harms, often (though not always) failing to find a privacy interest where a person has voluntarily shared information or is not in complete seclusion. [105]

In the context of government surveillance, as one of us has noted, **_robots_**, like other digital technology, raise the problem of the much maligned "third-party doctrine." [106] In Fourth Amendment jurisprudence, if less sensitive (that is, non-content) information is shared with a third party, no warrant is required. [107] In an age where nearly all information is shared with third parties, privacy protection vis-a-vis the government has been severely limited. There have been recent signals that the Supreme Court may move away from this approach, but it is currently still good **_law_**. [108] Because many **_robots_** will share information with third parties, the Fourth Amendment may not protect **_robot_** [*1001] owners from government surveillance. Because, however, protection for the home environment is so central to Fourth Amendment jurisprudence, home **_robots_** may spark a conflict between the third-party doctrine and historic protection for the home. [109]

Last but certainly not least, home **_robots_** pose a challenge for privacy **_law_** because of the **_law_**'s focus on one-time notice and consent. This focus ignores the dynamic nature of human-**_robot_** interactions. One-time notice upon purchase of a **_robot_** is not the same as dynamic feedback in the actual moment that a **_robot_** is observing through walls. [110] Moreover, when digital privacy problems no longer occur through a computer screen, notice becomes challenging. [111] What constitutes adequate notice of surveillance in a shared physical environment, what

---

[102] Hartzog, supra note 5, at 821-22.

[103] Meg Leta Jones, Privacy Without Screens & The Internet of Other People's Things, _51 Idaho L. Rev. 639, 640 (2015)._

[104] Danielle Keats Citron, The Privacy Policymaking of State Attorneys General, _92 Notre Dame L. Rev. 747, 749 (2016)._

[105] Danielle Keats Citron, Mainstreaming Privacy Torts, _98 Calif. L. Rev. 1805, 1826-28 (2010)._

[106] Kaminski, supra note 5.

[107] _Smith v. Maryland, 442 U.S. 735, 746 (1979); United States v. Miller, 425 U.S. 435, 446 (1976)._

[108] Kaminski, supra note 5, at 670; see also _Riley v. California, 134 S. Ct. 2473, 2490 (2014); United States v. Jones, 132 S. Ct. 945, 957 (2012)._

[109] Kaminski, supra note 5, at 669-70; see also Ferguson, supra note 3, at 840.

[110] Sedenberg, Chuang & Mulligan, supra note 51, at 584 (encouraging **_robot_** designers to embrace dynamic consent models).

[111] See, e.g., Christopher Wolff & Jules Polonetsky, An Updated Privacy Paradigm for the "Internet of Things" 4 (Nov. 19, 2013) (unpublished manuscript), _https://fpf.org/wp-content/uploads/Wolf-and-Polonetsky-An-Updated-Privacy-Paradigm-for-the-%E2%80%9CInternet-of-Things%E2%80%9D-11-19-2013.pdf_.

constitutes real consent, and whether such surveillance should be opt-in or opt-out are all challenging issues for privacy *law*. [112]

II. Technological Solutions

Updating privacy *law* is not easy. But privacy protections may be necessary for the uptake of new technologies like home *robots*. Some people might be aware of what they are bringing into their homes - and may not accept it. If there are several major privacy failures, people may throw a technology out. Technologists may want to design *robots* to mitigate privacy harms both because that approach is ethical, and because without responsible design, these technologies are unlikely to be widely accepted or adopted.

Regulators, too, can benefit from better understanding the role of technological design in mitigating or preventing privacy harms. By understanding what is possible through technological design, regulators can broaden a currently blunt toolkit in ways that may benefit both users and nascent technological fields. We focus here on the measures technologists can take, and in Part III below consider whether and how the *law* can encourage the adoption of these measures.

[*1002] The concept of building values into code or design has significant history. It has long been a principle in Internet *law* that people may be regulated not just by *law* but by code - that is, by technological design. [113] If technology significantly changes an environment, it might also be designed to mitigate the effects of those changes on social values. [114] Thus if *robots* effectively break down the walls of your home, either by walking around them or by seeing through them, *robots* might be designed to functionally reinstate those walls through other technological means.

The idea of looking to technological design to solve these problems is often referred to in the privacy context as "privacy by design." In 1997, Dr. Ann Cavoukin, the Information & Privacy Commissioner for Ontario, Canada, came up with principles for privacy by design. [115] Cavoukin proposed that privacy protections should be proactive, not reactive or remedial; privacy should be the default; privacy should be embedded into design; privacy should be seen as positive sum, not zero sum; designers should aim for end-to-end security; designers should aim for visibility and transparency; and designers should respect user privacy. [116] The German Federal Commissioner for Data Protection, Peter Schaar, later articulated six principles for privacy by design: data minimization; controllability: transparency; data confidentiality (security); data quality; and the possibility of segregation in multi-user environments. [117] In 2012, the FTC stated its reliance on privacy by design principles. [118]

---

[112] See, e.g., Fed. Trade Comm'n, Facing Facts: Best Practices for Common Uses of Facial Recognition Technologies 7, 12, 14-15 (2012), *https://www.ftc.gov/sites/default/files/documents/reports/facing-facts-best-practices-common-uses-facial-recognition-technologies/121022facialtechrpt.pdf*; *Jones, supra* note 103, at 645, 652-53; Margot E. Kaminski, When the Default Is No Penalty: Negotiating Privacy at the NTIA, 93 Denver L. Rev. 925, 931 (2016) (describing NTIA negotiations around facial recognition policy); Peppet, supra note 3, at 140, 146.

[113] Joel R. Reidenberg, Lex Informatica: The Formulation of Information Policy Rules Through Technology, *76 Tex. L. Rev. 553, 554 (1998)* ("Technological capabilities and system design choices impose rules on participants."); Lawrence Lessig, Code: And Other *Laws* of Cyberspace 6 (1999); Lawrence Lessig, Code: Version 2.0, at 5-6 (2006) [hereinafter Lessig, Code: Version 2.0].

[114] Harry Surden, Structural Rights in Privacy, *60 SMU L. Rev. 1605, 1606-08 (2007);* Paul Ohm & Jonathan Frankle, Proof of Work: Learning from Computer Scientific Approaches to Desirable Inefficiency (unpublished manuscript) (on file with the author).

[115] See David Krebs, "Privacy by Design": Nice-To-Have or a Necessary Principle of Data Protection *Law*?, 4 J. Intell. Prop. Info. Tech. & E-Commerce L. 2, 2 (2013).

[116] Ann Cavoukin, Privacy by Design: 7 Foundational Principles, *https://www.iab.org/wp-content/IAB-uploads/2011/03/fred_carter.pdf* (last visited May 17, 2017).

[117] Peter Schaar, Privacy by Design, 3 Identity Info. Soc'y 267, 273 (2010).

The principles of privacy by design are closely related to the Fair Information Practice Principles ("FIPPs") that form the foundation of many privacy _laws_ around the world. [119] The FIPPs were first proposed and named in **[*1003]** a 1973 report by an advisory committee in the Department of Health, Education, and Welfare. [120] The committee established a Code of Fair Information Practices that included the following requirements: no record-keeping systems may be secret; individuals must be able to find out what information is in a system and how it is used; individuals must be able to prevent information obtained for one purpose from being used for other purposes; individuals must be able to correct or amend records; and any organization keeping records must take precautions to prevent misuse and must assure the reliability of the data for their intended use. [121] A later commission in 1977 expanded the five HEW principles into eight principles. [122] And in 1980, the Organization for Economic Cooperation and Development ("OECD") adopted eight principles in its Guidelines on the Protection of Privacy and Transborder Flows of Personal Data. [123] The eight FIPPs outlined by the OECD form the foundation for privacy _laws_ around the world and are discussed in greater detail in Part II.A below.

Apart from attempts at establishing broad principles, privacy by design has arisen in specific technological applications. For example, in the area of ubiquitous computing, or ubicomp, discussions of privacy often entail implementing privacy protections through technological design. Ubiquitous computing has been defined as "making many computers available throughout the physical environment, while making them effectively invisible to the user." [124] As early as 1993, ubicomp researchers proposed "design for privacy" principles, including specific design suggestions. [125] One researcher proposed six principles: notice, choice and consent, proximity and locality, anonymity and pseudonymity, security, and access and recourse. [126] These principles largely reflect the FIPPs. Another set of researchers proposed the idea of "situational faces," whereby users could create appropriate "faces" or **[*1004]** user profiles for particular environments. [127] Those researchers proposed that designers also notify

---

[118] Edith Ramirez, Fed. Trade Comm'n Commissioner, Remarks at Privacy by Design Conference in Hong Kong 1 (June 13, 2013), _https://www.ftc.gov/sites/default/files/documents/public_statements/privacy-design-and-new-privacy-framework-u.s_. federal-trade-commission/120613privacydesign.pdf.

[119] See Dep't Homeland Sec., DHS Mem. No. 2008-1, Privacy Policy Guidance Memorandum (2008), _https://www.dhs.gov/xlibrary/assets/privacy/privacy_policyguide_2008-01.pdf_; OECD Privacy Principles, OECD Privacy,_http://oecdprivacy.org/_ (last visited May 17, 2017).

[120] Advisory Comm. on Automated Personal Data Systems, Dep't of Health, Education and Welfare, Records, Computers, and the Rights of Citizens (1973), _https://www.justice.gov/opcl/docs/rec-com-rights.pdf_.

[121] Id. at xxiii-xxxv, 40-41.

[122] Privacy Protection Study Comm'n, Protecting Privacy in an Information Society (1977), _https://epic.org/privacy/ppsc1977report/_.

[123] OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, OECD [hereinafter OECD Guidelines], _http://www.oecd.org/internet/interneteconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm_ (last visited May 17, 2017).

[124] Mark Weiser, Some Computer Science Issues in Ubiquitous Computing, 36 Comm. of the ACM, no. 7, July 1993, at 75, 75.

[125] Victoria Bellotti & Abigail Sellen, Design for Privacy in Ubiquitous Computing Environments, in Proceedings of the Third European Conference on Computer-Supported Cooperative Work 77, 77 (Giorgio De Michelis et al. eds., 1993) (ebook).

[126] Marc Langheinrich, Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems, in Ubicomp 2001: Ubiquitous Computing 273, 273 (Gregory D. Abowd et al. eds., 2001).

[127] Scott Lederer et al., A Conceptual Model and a Metaphor of Everyday Privacy in Ubiquitous Computing Environments (2002), _http://digitalassets.lib.berkeley.edu/techreports/ucb/text/CSD-02-1188.pdf_.

users at the boundaries between different information environments. [128] The Internet of Things ("IoT") raises similar issues to ubiquitous computing, and has prompted similar discussions about privacy and design. [129]

In this Part, we turn to technological design and how it might mitigate the above-named privacy harms. **_Robot_** privacy is a subset of the field of Human-**_Robot_** Interaction ("HRI"). [130] Below we build on our identification of **_robot_**-related privacy harms to propose a set of privacy principles, identify existing technological solutions, and pinpoint the more challenging technological problems that lie ahead.

A. Principles

We posit above that home **_robots_** will raise three types of privacy problems: (1) data privacy problems; (2) boundary management problems; and (3) social/relational problems. The FIPPs are a useful resource for addressing data privacy problems across many kinds of technologies, including **_robots_**. [131] Data privacy problems - the privacy problems raised by the collection and maintenance of vast systems of records on individuals - are similar across technologies. It is thus unsurprising that the FIPPs would be applicable here. But the FIPPs, as currently practiced, do not adequately address the second and third types of privacy harms we have identified: boundary management problems and social/relational problems.

Thus, we begin with a selection of principles from the FIPPs, but propose two additional principles for technologists to follow and perhaps for regulators to enforce. Roboticists should design home **_robots_** with an eye to the FIPPs principles of data minimization, purpose specifications, and use limitations, discussed at greater length below. To this list, we add two additional principles: honest anthropomorphism, and dynamic feedback and participation. We caution that the incorporation of these principles into design should be an integral part of the overall design process, rather than a post **[*1005]** facto afterthought. [132] And we believe it important to concretize these principles with examples of actual technologies that could provide solutions, which we do below.

To the extent that the privacy problems raised by home **_robots_** involve the collection and analysis of large quantities of information, the FIPPs are applicable. The eight FIPPs, as articulated by the OECD and incorporated into privacy **_laws_** around the world, are: collection limitation (also known as data minimization), purpose specification, use limitation, data quality, security safeguards, openness, individual participation, and accountability. [133] The collection limitation principle (or data minimization principle) states, "there should be limits on the collection of personal data, and such [collection] should be … [done], where appropriate, with the knowledge or consent of the data subject. [134] This principle pushes back on any practice of gathering information indiscriminately, without a purpose in mind and without knowledge and consent of the data subject.

The purpose specification principle relates to collection limitation, in that it cautions against indiscriminate collection and use, and requires data collectors to specify the purpose of data collection at the time of collection. (One problem for the FIPPs is a lack of indication of just how broad or narrow that purpose may be. We discuss this further below.) The use limitation principle then limits subsequent use of the data to the fulfillment of those stated purposes - or for other purposes only with the consent of the data subject.

---

[128] Id. (referred to as the Boundary Principle).

[129] See generally Luigi Atzori et al., The Internet of Things: A Survey, 54 Computer Networks 2787 (2010).

[130] Rueben & Smart, supra note 23, at 2.

[131] Sedenberg, Chuang & Mulligan, supra note 51, at 580 (applying a version of the FIPPs to information privacy harms posed by therapeutic **_robots_**).

[132] See Cavoukin, supra note 116.

[133] OECD Guidelines, supra note 123.

[134] Id.

The envisioned process dictated by these three principles entails stating a collection purpose, notifying and obtaining consent from the data subject, gathering only enough information to fulfill that purpose, and using the information to fulfill that purpose - or returning to the data subject for consent to use it for other purposes. [135]

While the FIPPs are directed towards regulators that design the **_laws_** that govern databases, they can readily be understood from the perspective of designers as well. The three principles of data minimization, purpose specification, and use limitation together require technologists to know and state why they are gathering information before they begin to gather it, and restrict the use of information to those purposes. They require technologists to notify and usually obtain consent from data subjects upon information collection, and again upon broader use beyond the expected purposes. These are solid core principles for **_robot_** design, though we discuss their limitations below.

 [*1006] The remaining FIPPs are more focused on the governance of data once it is in a database, and are thus less applicable to the **_robot_**-as-interface design questions we discuss here. Data quality, security safeguards, openness, accountability, and individual participation (which, in brief, is defined largely as a right to obtain and challenge information held in the database) are important - as is the requirement that robotics companies create data deletion policies so as not to indefinitely retain information for no good reason - but these largely address concerns outside the scope of this project. These FIPPs are more relevant to discussions of informational due process [136] : the ability of individuals to understand and challenge decisions made about them based on big data analytics. Our focus here is more on the design of home **_robots_** as an information-gathering interface with humans.

We thus focus our efforts here on data minimization, purpose specifications, and use limitations. Home **_robots_** should gather information only for a specific, articulated purpose or purposes, [137] should attempt to limit the information gathered to information necessary for that purpose, should avoid gathering sensitive information, and should share information no more than is necessary for the stated purpose.

To normatively ground the process of designating purpose limitations, we suggest technologists consider Nissenbaum's concept of contextual integrity. [138] A **_robot_**'s stated purpose should be connected to the context in which the user understands the **_robot_** to be operating, and use of the data gathered by that **_robot_** should be limited to that context. For example, a vacuuming **_robot_** like the Roomba might appropriately have the stated purpose of effectively and/or efficiently cleaning a room. It should not have the broader purpose of, say, making medical diagnoses about a user. The use of information gathered by a vacuuming **_robot_** in the course of cleaning a room to infer broader behavioral information about a user would violate that user's sense of the **_robot_**'s purpose in context. Thus, contextual integrity can serve as a helpful guide for establishing purposes in line with a user's actual expectations of the technology and the information environment it inhabits.

The three FIPPs of data minimization, purpose specification, and use limitation provide helpful starting guidelines for technologists and regulators. **[*1007]** They do not, however, adequately address our second and third types of privacy harms. They fail to fully articulate how a **_robot_** should be designed so as to adequately mitigate boundary-management problems, or how a **_robot_** should be designed to address the privacy harms caused by the **_robot_** being treated as a social actor.

---

[135] Id.

[136] See Danielle Keats Citron & Frank Pasquale, The Scored Society: Due Process for Automated Predictions, _89 Wash. L. Rev. 1, 27-28 (2014);_ see also Sedenberg, Chuang & Mulligan, supra note 51, at 580-81 (largely focusing on the ability of the users of therapeutic **_robots_** to access and amend information).

[137] There is some discussion between the co-authors of this Article about what constitutes "gathering" - whether it includes only collection that results in storage and/or dissemination, or whether it also extends to the collection of information stored in volatile memory but never committed to disk or sent over the network boundary of the **_robot_**. This conversation about whether gathering without meaningful storage or dissemination "counts" as a privacy harm is a familiar one. We note but do not resolve this discussion here.

[138] Nissenbaum, supra note 71.

First, the FIPPs do not fully mitigate the **_robot_**-specific boundary-management problems we identify: that **_robots_** can sense through or move around physical boundaries humans use to manage privacy. The data minimization principle may in practice counsel that **_robots_** should not gather particularly sensitive information, and that information may be kept behind walls or behind physical barriers. But the data minimization principle does not address the problem that humans may underestimate or misunderstand **_robots_**' sensory or physical capacities.

We thus suggest a more dynamic type of notice, which we refer to as a principle of "dynamic feedback and participation." [139] Under the principle of dynamic feedback and participation, technologists should design **_robots_** to regularly indicate to users how their presence changes an information environment, including by indicating when physical and sensory barriers are not in fact barriers to a **_robot_**.

This is not necessarily a one-way flow of information. **_Robot_** designers could design **_robots_** to pick up social feedback from users, and could either relay that feedback to **_robot_** operators or companies, or incorporate it directly into how a **_robot_** operates. Perhaps **_robots_** could be designed to detect or be alerted to signs of "privacy outrage" - when a user is particularly offended by a perceived privacy intrusion. A **_robot_** could in real time adjust its behavior accordingly or notify their companies that something needs to be fixed.

This differs significantly from the notice contemplated in the FIPPs. In framing privacy issues for regulators, the FIPPs focus on more static, one-time notice and consent upon gathering or distribution. [140] This fails to encourage designers to use more effective forms of notice, built-in to a technology and recurring throughout a user's interactions with it. [141] Thus, the principle of dynamic feedback that we propose here is not a one-time notice to a **_robot_** user, but a designed process of notifying a user of what a **_robot_** is actually doing, and trying to incorporate a user's response into how the **_robot_** treats that information environment.

Our second non-FIPPs principle of "honest anthropomorphism" is aimed at addressing the privacy problems raised by the fact that a **_robot_** can **[*1008]** be designed to be treated as a social actor. The principle of honest anthropomorphism is as follows: **_Robot_** designers should not use anthropomorphism to deliberately mislead users as to privacy practices. If anything, roboticists should explore using anthropomorphic features to provide better notice to users of what a **_robot_** is actually doing. [142]

Our five principles for privacy-sensitive **_robot_** design thus are: data minimization, purpose specification, use limitation, dynamic feedback, and honest anthropomorphism. Implementing these principles will not always be easy or obvious. For example, dynamic feedback is a significant design challenge. Notifying users on a computer screen is in some ways easier than notifying them of surveillance in the real world, because there is not always a clear moment that delineates the beginning of user interaction with a **_robot_** or IoT device, and there is no screen to post the notice on or make a user click through. Some forms of notice are far more effective than others, and designers should consider this. [143] When and how often to provide notice - initially, repeatedly, at certain times of day or at certain changes to the information environment - is also a challenging question. It may make sense to incorporate the Boundary Principle idea from ubiquitous computing, which urges technologists to notify users at least when an information environment has meaningfully changed. [144] It may also make sense for designers to consider ways of giving users the option of different "situational faces," protecting privacy to different degrees depending on time of day, location within the home, or social setting (such as, "dinner with friends" versus "time with loved ones").

---

[139] This is related to the idea of dynamic consent models proposed in Sedenberg, Chuang, & Mulligan, supra note 51, at 584.

[140] Woodrow Hartzog, The Inadequate, Invaluable Fair Information Practices, _76 Md. L. Rev. 952, 974, 975 (2017)._

[141] _Id. at 979._

[142] Rueben & Smart, supra note 23, at 31.

[143] Rebecca Balebako et al., Is Your Inseam a Biometric? Evaluating the Understandability of Mobile Privacy Notice Categories, Carnegie Mellon U. CyLab (2013), _https://www.cylab.cmu.edu/files/pdfs/tech_reports/CMUCyLab13011.pdf_.

[144] See Lederer et al., supra note 127.

Purpose specification is also particularly challenging in the age of big data. The purpose specification principle, again, suggests that information be gathered only if it advances a particular purpose or use. If every piece of information is useful - for data mining, machine learning, behavioral advertising, communication, or navigation - then limiting information gathering to a particular purpose will not do much to protect privacy. In other words, some **_robots_** may have a narrow purpose; thus, their design may meaningfully benefit from application of purpose specification. But other **_robots_** with broader purposes may have such broad mandates that every piece of information is plausibly "useful." This is a problem not just for **_robots_**, but also for big data writ large. We return to our above suggestion that technologists designate **_robot_** purposes with an understanding of how users perceive the **_robot_**'s role, and with the goal of preserving contextual integrity.

 [*1009] Designers can still think about whether all functions are necessary for every kind of **_robot_**. A talking toy may need to audio record and store audio recordings so it can learn and respond, but a cleaning **_robot_** probably does not need to do so. A social **_robot_** may need to employ facial recognition, but a **_robot_** built to clean your gutters or deliver you snacks probably will not. Designers can learn to think in principled ways about why they include certain technologies, how long (and how securely) information really needs to be stored, and with whom it needs to be shared.

For example, the makers of the Roomba 980 may want to contemplate whether the amount of data gathered for navigation is truly necessary for the Roomba's purpose: to clean the floor. As discussed, it may be the case that the Roomba's core use for floor cleaning is in fact as well or better served by a less data-intensive navigation technology. A Roomba owner might be notified of what information the Roomba gathers, and not just at the initial purchase. A Roomba owner might be given the option to put in virtual walls, not just to prevent objects from being run over, but also to guard private areas. And perhaps non-owners impacted by the Roomba should also be given privacy choices.

Not all of our five principles will be addressed through technological design. Use limitation, for example, is really more about a company practice than it is about the design of an interface. Purpose specification similarly and significantly involves company practices, though it also can involve designing a technology around a particular use or purpose. But implementing data minimization, dynamic feedback, and honest anthropomorphism will involve significant design choices. For this reason, we next turn to the technologies available for implementing these principles.

B. Technical Solutions Using Existing Technology

 Principles are useful as a baseline, but technologists and regulators will both be aided by concrete examples of the available toolkit for implementing them. There are various technologies that can be used to minimize data gathering and encourage dynamic participation in determining and customizing a home **_robot_**'s privacy settings. Some of these technologies could easily be incorporated into existing systems; others will require more work to adapt a technology for a **_robot_**-specific application.

First, users could be given the option of constraining **_robot_** navigation - that is, preventing **_robots_** from entering certain spaces or interacting with certain objects. This would both minimize the data collected, including particularly sensitive data, and reinstate certain physical barriers in the home through technological means. One big caveat to this approach is that such **[*1010]** constraints could themselves communicate information, such as which areas are considered to be most important or most private in a house. [145]

Private areas or private objects could be designated as obstacles that **_robots_** avoid, using motion-planning algorithms. [146] Designers could add a temporal dimension to a **_robot_**'s map, designating a particular space as an

---

[145] Rueben & Smart, supra note 23, at 18; Calo, supra note 32, at 198 ("The way we use human-like **_robots_** will be fixed in a file. Suddenly our appliance settings will not only matter, they also will reveal information about us that a psychotherapist might envy.").

[146] Steven M. LaValle, Planning Algorithms 105-52 (2006).

obstacle or an open area, depending on the time of day. [147] This could allow users to keep a ***robot*** out of rooms at some times, while allowing them in at others. Users might be given the option, via a graphical interface, to indicate private objects, regions, or time periods. [148] Researchers have studied the use of markers to indicate private areas, versus hand gestures. [149] And the principle of feedback could be incorporated too, by using haptic feedback to alert remote operators of ***robots*** when a ***robot*** nears a restricted area. [150]

Another way to achieve both data minimization and individual feedback through constraints on navigation is to use semantic mapping. A semantic map adds higher-level conceptual information to a map, beyond mere metric measurements, including labels for persons, places, and things. [151] These meaning-filled labels could then inform a ***robot***'s decisions. A human could assign these labels ("bedroom, do not enter") or a ***robot*** could actually infer them if given the right information ("if bed, then bedroom, then do not enter"). [152] The former would likely be more accurate and more effective, but the latter could afford even those users who fail to assign specific labels to specific places some level of privacy protection.

 [*1011] Constraining ***robot*** navigation with respect to proximity to people has been a central problem of HRI. Researchers have studied what kinds of approach behaviors make humans most uncomfortable. [153] Numerous ***robot*** behaviors have consequently been created with the goal of preserving personal space. [154] With respect to mapping discussed above, it will be difficult but worth exploring how to program a ***robot*** to track and obscure mobile objects - such as the personal space around a particular person. [155]

Interestingly, studies of personal distance have addressed other aspects of ***robot*** design, including eye contact [156] ; whether ***robots*** were designed with legs or wheels [157] ; and how speed, movements, and even headlight brightness can be scaled based on proximity to increase participant comfort. [158]

---

[147] Rueben & Smart, supra note 23, at 23.

[148] Id. at 32-33.

[149] Nisarg Raval et al., MarkIt: Privacy Markers for Protecting Visual Secrets, 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing; Matthew Rueben et al., Evaluation of Physical Marker Interfaces for Protecting Visual Privacy from Mobile ***Robots***, 2016 IEEE Int'l Symp. on ***Robot*** & Human Interactive Comm. 787, 787-88, 793.

[150] Frederik Ryden & Howard Jay Chizeck, A Method for Constraint-Based Six Degree-of-Freedom Haptic Interaction with Streaming Point Clouds, 2013 IEEE Int'l Conf. on Robotics & Automation 2353, 2353; Rueben & Smart, supra note 23, at 25.

[151] Cipriano Galindo et al., Multi-Hierarchical Semantic Maps for Mobile Robotics, 2005 IEEE IRS/RSJ Int'l Conf. on Intelligent ***Robots*** & Systems 3492; Rueben & Smart, supra note 23, at 24.

[152] Rueben & Smart, supra note 23, at 24.

[153] See John Travis Butler & Arvin Agah, Psychological Effects of Behavior Patterns of a Mobile Personal ***Robot***, 10 Autonomous ***Robots*** 185 (2001); Leila Takayama & Caroline Pantofaru, Influences on Proxemic Behaviors in Human-***Robot*** Interaction, 2009 IEEE/RSJ Int'l Conf. on Intelligent ***Robots*** & Systems 5495; Michiel Joosse et al., BEHAVE-II: The Revised Set of Measures to Assess Users' Attitudinal and Behavioral Responses to a Social ***Robot***, 5 Int'l J. Soc. Robotics 379 (2013).

[154] Rueben & Smart, supra note 23, at 24.

[155] Id. at 23.

[156] Jonathan Mumm & Bilge Mutlu, Human-***Robot*** Proxemics: Physical and Psychological Distancing in Human-***Robot*** Interaction, 2011 ACM/IEEE Int'l Conf. on Human-***Robot*** Interaction 331, 336 (noting that people who dislike ***robots*** maintained a greater physical distance when a ***robot*** was looking at them).

[157] Sandra Y. Okita et al., Captain May I? Proxemics Study Examining Factors that Influence Distance Between Humanoid ***Robots***, Children, and Adults, During Human-***Robot*** Interaction, 2012 ACM-IEEE Int'l Conf. on Human-***Robot*** Interaction 203.

[158] Zachary Henkel et al., Evaluation of Proxemic Scaling Functions for Social Robotics, 44 IEEE Transactions on Human-Machine Systems 374, 375, 377 (2014).

One worry, discussed above, is that in calibrating **_robot_** design features to make people more comfortable, **_robot_** designers may give users a false sense of comfort. [159] Users may be concerned about their lack of knowledge of a **_robot_**'s actual abilities - for example, researchers studying a social **_robot_** in the workplace could not tell its sensing capabilities by its appearance, and expressed a desire to be better notified. [160] We thus return to our principle of honest anthropomorphism to suggest that such design elements be coupled with comparable constraints on **_robot_** navigation or surveillance, or at least with additional notification systems. If a **_robot_** is designed to make its user feel more comfortable by lowering its eyes and avoiding eye contact, it should not simultaneously be recording its user from a neck-mounted camera, for example, without some additional alert.

 [*1012] A second approach to data minimization is to constrain **_robot_** perception, rather than, or in addition to, navigation. [161] Given **_robots_**' reliance on environmental information for much of their functioning, one recurrent concern is that there may be a significant tradeoff between utility (including functional navigation) and visual privacy. [162] Several researchers, however, have evaluated this privacy-utility tradeoff and found it feasible for **_robots_** to complete tasks with effective filters in place. [163] One study looked more specifically at different methods of constraining perception and ranked them by utility; we return to this study and its outcome below. [164]

There are multiple methods of constraining perception. First, one could use methods of navigation that do not involve visual information. For example, **_robots_** can navigate using a depth camera instead of a color camera. [165] The depth camera, however, has its limitations and only works within a certain range. [166] Additionally, a depth camera is still a camera, even if it does not contain color information. [167]

For **_robots_** that use cameras, there are a range of ways of post-processing images to protect privacy. You can reduce the resolution of an image, by pixelating it. [168] You can smooth the image by allowing pixels to influence the values of neighboring pixels, by blurring it. [169] You can remove pixels and place a black box over the objectionable part of an image, by redacting it. [170] Or you can remove pixels and replace them with what is behind an offending

---

[159] Rueben & Smart, supra note 23, at 27-28.

[160] Min Kyung Lee et al., Understanding Users' Perception of Privacy in Human-**_Robot_** Interaction, 2011 ACM/IEEE Int'l Conf. on Human-**_Robot_** Interaction 181, 182.

[161] Rueben & Smart, supra note 23, at 18.

[162] Id. at 25.

[163] Alexander Hubers et al., Video Manipulation Techniques for the Protection of Privacy in Remote Presence Systems, 2015 ACM/IEEE Int'l Conf. on Human-**_Robot_** Interaction Extended Abstracts 59-60; Rueben & Smart, supra note 23, at 25.

[164] Adam Erdelyi et al., Adaptive Cartooning for Privacy Protection in Camera Networks, 11 IEEE Int'l Conf. on Advanced Video & Signal-Based Surveillance 44 (2014).

[165] Chenyang Zhang et al., Privacy Preserving Automatic Fall Detection for Elderly Using RGBD Cameras, in Computers Helping People with Special Needs 625 (Klaus Misenberger et al. eds., 2012).

[166] Id.

[167] See, for example, the discussion of TSA body scanners. Yofi Tirosh & Michael Birnhack, Naked in Front of the Machine: Does Airport Scanning Violate Privacy?, _74 Ohio St. L.J. 1263, 1263-64 (2013);_ Lisa Brownlee, Growing List of Privacy Advocates Condemns TSA's New Body Scan Policy, Forbes (last updated Jan. 16, 2016, 10:05 AM), _http://www.forbes.com/sites/lisabrownlee/2016/01/14/growing-list-joins-tsa-body-scan-fight/_; Security Executives, Full-Body Scanning - A Search for Balance Between Privacy and Security, Homeland Security (Feb. 9, 2016), _https://medium.com/homeland-security/full-body-scanning-a-search-for-balance-between-privacy-and-security-1e533d4870c8#.2h1m4wyd5_.

[168] Rueben & Smart, supra note 23, at 19.

[169] Id.

[170] Id.

object. The last option, replacement, is also known as "inpainting" or "image completion." [171] Of these techniques (excluding replacement), one **[*1013]** study found privacy most preserved by redaction, then pixelation, and least by blurring. [172] The above-referenced study calculating utility versus privacy found that abstracting the image provided the most utility but the least privacy, followed by blurring; and pixelation provided the most privacy but least utility. [173]

There are additional techniques for post-processing an image. You can employ brush-stroke effects (painting techniques) that incidentally remove identifying details. [174] You can manipulate the focus in an image, which causes people to ignore the parts of the image that are out of focus. [175] You can remove colors and textures, and represent 3-D models of objects with surface contours.

A particular challenge for home **_robots_** is what to do about personally identifiable information belonging to **_robot_** owners. To mitigate the intrusion caused by facial recognition technology, designers can morph faces so that they are unrecognizable. [176] One study explored whether de-identification using pixelation, edge detection, and abstractions could provide greater privacy. [177] The study found that even with these protections, people could still be identified by shirt color. [178] There are, unfortunately, multiple ways to identify people beyond their faces: clothes, gait, behavior, and when and where they enter an environment, particularly when people are repeat actors in a particular space. [179]

These existing technologies can be deployed or adapted to implement our five privacy principles. By constraining **_robot_** navigation or perception, or providing additional processing to gathered information, technologists can make significant contributions to enabling data minimization. This would prevent the gathering of information out of context and unrelated to a **_robot_**'s perceived purpose. For users to have meaningful input into the process, designers will have to build interfaces. One of us has researched three different interfaces for specifying visual privacy preferences to a **_robot_**. The usability **[*1014]** of these interfaces was found to depend on the scenario. Building user interfaces that allow users to both know and influence what information a **_robot_** sees, gathers, and uses would significantly implement our principle of dynamic feedback.

C. Technical Solutions That Require Research and Development

Some privacy concerns remain difficult to address using current technology. This Section discusses problems that technology cannot reliably solve, at least not yet. We start by discussing some limits on our earlier discussion of imposing constraints on **_robot_** navigation or perception. We then turn to the challenges inherent in trying to operationalize privacy settings, whether they are set by individual users or by a manufacturer for all users.

Current technology faces significant limitations in reasoning from privacy settings, in detecting sensitive objects or scenarios, and in understanding context. **_Robots_** are good at rules, but bad at making analogies the way a human

---

[171] Id. at 20.

[172] Pavel Korshunov et al., Crowdsourcing Approach for Evaluation of Privacy Filters in Video Surveillance, 2012 ACM Workshop on Crowdsourcing for Multimedia 35, 39.

[173] Erdelyi, supra note 164.

[174] Jingwan Lu et al., Interactive Painterly Stylization of Images, Videos and 3D Animations, 2010 ACM SIGGRAPH Symp. on Interactive 3D Graphics & Games 127, 132-33.

[175] Forrester Cole et al., Directing Gaze in 3D Models with Stylized Focus, 2006 Eurographics Symp. on Rendering 377, 385 .

[176] Pavel Korshunov & Touradj Ebrahimi, Using Face Morphing to Protect Privacy, 2013 IEEE Int'l Conf. on Advanced Video & Signal Based Surveillance 208.

[177] Qiang Alex Zhao & John T. Stasko, Evaluating Image Filtering Based Techniques in Media Space Applications, 1998 ACM Conf. on Computer Supported Cooperative Work 11, 17.

[178] Id.

[179] Rueben & Smart, supra note 23, at 22.

would, or understanding context as a human might. This makes operationalizing a particular user's privacy settings - or trying to establish more general privacy settings based on shared norms - a significant challenge in practice.

We then turn to notice problems: challenges around letting users know what a **_robot_** is in fact doing. We close this Section by asserting that even with known technologies, to properly address privacy problems we must engage in whole-system testing, applying and testing technologies with real **_robots_** in real-world settings. Throughout, we explain how these technological challenges implicate our five principles of data minimization, purpose specifications, use limitations, honest anthropomorphism, and dynamic feedback and participation.

1. Constraining **_Robot_** Navigation and Perception

We begin by addressing some technical limitations on the above-discussed constraints on **_robot_** navigation and perception. These solutions, as noted, primarily go to the principle of data minimization: trying to minimize the amount of information a **_robot_** gathers while still leaving the **_robot_** functional.

Currently, with respect to constraining **_robot_** perception, it is difficult to perform convincing object replacement in a live video stream. [180] It is hard to make the video run fast enough while maintaining a map of what is behind **[*1015]** the object so as not to tip off to the viewer that something is missing from the image. This problem may be solvable, however, in the medium term.

More fundamentally, research should focus on how to navigate using privacy-preserving sensors. To protect privacy, we might want to collect only a particular kind of sensor data, or less sensor data. **_Robots_** can gather information in a variety of ways; cameras gather color image data, whereas lasers gather distance data that can be used to approximately reconstruct 3-D surfaces. Navigation algorithms need to be developed to use different kinds and different amounts of sensor information and still work reasonably well.

Another fundamental problem with constraining **_robot_** navigation is that the privacy protection methods discussed above have the potential to call attention to private objects, regions, or people. Onlookers, operators, or malicious actors could use the avoidance of particular objects or rooms to infer both value and location, and consequently overcome efforts to obscure private objects. [181]

2. Reasoning from and Implementing Privacy Settings

A second class of challenges arises around the process of reasoning from privacy settings, whether they are set by a particular user or by a company attempting to protect multiple or all users. Setting privacy preferences and making context-specific conclusions from them will present difficulties. These difficulties largely implicate our principle of dynamic feedback and participation, and implicate use limitations when a company attempts to limit sharing of information out of context.

In general, **_robots_** are good at following rules. However, these rules need to be specified precisely, in terms that the **_robot_** can reliably measure with its sensors. Since our privacy concerns are usually highly contextual and depend on subtleties, this makes them hard to articulate to a **_robot_**. **_Robots_** are bad at drawing analogies and conclusions in the way that people do. And they are currently bad at detecting context.

When it comes to reasoning from a set of privacy settings, it would be useful to be able to indicate privacy settings not just on an object-by-object basis but by object type. For example, an owner might want to label as private "all my documents," or more specifically "all my legal documents," rather than having to tell a **_robot_** to ignore each specifically identified document. There are significant technical challenges to this semantic labeling, in making it usable and not too burdensome for most **_robot_** owners. It might be feasible for a **_robot_** to detect a document, by looking for a piece of paper with writing on it. However, interpreting it as a "legal document" requires more work. If a

---

[180] For early work on technology that makes this at least partly possible, see Paul E. Debevec et al., Modeling and Rendering Architecture from Photographs: A Hybrid Geometry-and Image-Based Approach, 23 Int'l Conf. on Computer Graphics and Interactive Techniques 11 (2006).

[181] Rueben & Smart, supra note 23, at 32.

"legal document" is defined as a document with legal words in it **[*1016]** (tort, litigation, etc.), then the **_robot_** must be provided an exhaustive list of these words, and some criteria for classification: does a "legal document" contain only one of these words, or 200 of them?

The process of semantic labeling can easily get bogged down in having to be extremely specific about things humans intuitively understand. Imposing the burden of this specificity on **_robot_** owners and users will make them unlikely to create detailed privacy settings. Thus, the challenge with semantic labeling is balancing (1) not creating too burdensome of a labeling task for users with (2) covering enough variations in types of objects or locations that important distinctions are not glossed over.

A second significant challenge for implementing privacy settings is detection. It is hard for **_robots_** to detect a type of object or situation identified in a privacy setting. Take the example of credit card information. If a privacy setting requires blurring out credit cards, the **_robot_** system needs to be able to recognize credit cards quickly and accurately enough that they can be blurred in every frame of the **_robot_**'s sensor feed without fail. Object detection in a constrained setting is a more-or-less solved problem. Object detection "in the wild" is a more difficult work in progress, but it is solvable in the medium term for enough kinds of objects to make it useful. The technology is fast improving, especially with the use of convolutional neural networks ("CNNs") that have been trained on large databases of labeled images. [182] But privacy protection is a demanding domain: whereas in most applications of object detection it is okay to just get it right most of the time, a single failure to detect a private object could cause a user serious harm.

Some detection problems remain truly challenging, however. First, detecting moods or emotions rather than objects is challenging. Second, even some kinds of object detection remain meaningfully hard. Third, even where accurate object detection is possible, distinguishing between categories of similar objects is not necessarily easy. And it is hard for a **_robot_** more generally to determine the norms of a particular situation.

Robots will struggle to detect things about humans like emotion, mood, and social cues. It can be hard for a **_robot_** to even detect the direction someone is looking. Thus, trying to design a **_robot_** to detect "privacy outrage" - when a user or owner needs more personal space or more alone time - will be difficult. Ideally, a user or owner should be able to establish high-level privacy settings, such as "do not enter when we are fighting." But detecting the relevant information for determining these kinds of settings - scene understanding, emotion detection, and as we discuss more below, context - will be challenging.

 [*1017] Even object detection still has considerable challenges in the privacy context, despite recent advances in the technology. Implementing privacy settings demands high recognition accuracy, where the **_robot_** either perfectly detects the objectionable information or knows when it is unable to be perfect so it can shut off a sensor feed. Take the example of a partially clothed person. A user might not want a **_robot_** to record him partially undressed. The first step in addressing this preference is to detect when the user is, in fact, unclothed. Nudity detectors could detect and mask the appropriate regions. [183] But these kinds of detectors face problems. To program a computer vision system on a **_robot_** to detect something, you again must first specify precisely what that something looks like: its color, perhaps its shape, size, and other characteristics. Saying "red" is not good enough, you have to give a range of pixel values that you consider to qualify as "red." This is hard for most objects, since their apparent color changes under different lighting conditions, but it is even harder for things like skin. Skin comes in all sorts of colors, and humans use other information to figure out where it is. Skin is usually wrapped around a human, which makes it easier to identify. However, this concept of "wrapped around a human" is hard to articulate to a computer vision system.

The more difficult underlying problem is "scene understanding." It is hard for a computer to perceive what is going on in an image, especially with no a priori knowledge of the context. This creates an object detection problem for objects that look similar but have very different social meanings. For your undressed user, a swimsuit, underwear, and lingerie are three types of clothing that might look very similar to a **_robot_**. Even if the **_robot_** detects that a

---

[182] See, e.g., Image-Net, _http://image-net.org/_ (last visited May 17, 2017).

[183] Margaret M. Fleck et al., Finding Naked People, 4 Eur. Conf. on Computer Vision 593 (1996).

person is only partially dressed, it may not be able to decide what this information means with respect to its marching orders.

Thus, even if we could reliably detect the things that we want to (skin, humans), many of our concerns about objects are wrapped in a social context. To automatically enforce this, we have to define, in very precise terms, using only the sensor data a _**robot**_ can gather, what is meant by words like "public" and "private," for example. Doing this is currently beyond the state of the art in computer vision, and is likely to remain so for some time, despite recent progress in more constrained domains (such as labeling images on the Web).

The broader problem behind implementing privacy settings is the fact that _**robots**_ are challenged by context-specific factors. This implicates both our principles of data minimization and use limitation. _**Robots**_ that operate beyond just one context will struggle with implementing Nissenbaum's conception of privacy as "contextual integrity." [184]

 [*1018] For example, consider a _**robot**_ that cleans office cubicles, while attempting to respect workers' privacy. It would be hard to write down a set of concrete rules (for example, "do not enter the cubicle from 9-10 a.m. (my weekly meeting time)" or "do not enter when there are two or more people inside") that will perfectly capture a preference that the _**robot**_ should stay away when a worker is in the middle of an important meeting. The _**robot**_ should ideally be able to figure out when a worker is in "an important meeting" - which is a "scene understanding" problem, discussed above. Scene understanding is made more difficult by the Frame Problem, which states that in order to determine the frame or context of a situation, you need to interpret the facts you are given about it. But in order to interpret the facts about a particular context, a _**robot**_ needs to know about the context, leaving it stuck.

Similarly, it might be useful for a _**robot**_ to detect a particular person's territory or belongings, and follow social conventions for respecting that territory. This would include detecting the relevant properties of potential territory markers, such as to whom the property belongs. Again, this is a difficult task because of the scene-level understanding required. For example, if I hand my umbrella to a friend, the _**robot**_ would need to detect whether the friend is borrowing my umbrella or if it was a gift, making it my friend's umbrella now. Otherwise, I would have to tell the _**robot**_ every time I gain or lose belongings. Other social cues indicating territory might also need to be detected by the _**robot**_, from a person's mere presence to how they are sitting.

In light of the significant challenges with detecting objects and reasoning about and from context, it may be wise to give up on perfect privacy filters. A probabilistic framework could instead provide filtering based on a user's desired confidence level. [185] If a _**robot**_ is unsure about whether a certain object or scenario is present in its sensor filter, a _**robot**_ belonging to a user with a lower desired confidence setting could take a risk and not filter anything. If the unsure _**robot**_ belongs to a user with a higher confidence setting, it could instead blur its entire feed to make sure any possibly private object is obscured.

3. Notice and Feedback Problems

 A third category of technical challenges involves creating notice and feedback systems for the privacy-conscious user. This set of problems implicates two of our principles: honest anthropomorphism, and dynamic feedback and participation.

It is hard to participate in _**robot**_ privacy settings if a user does not know what the _**robot**_ is actually doing. Anthropomorphic _**robots**_ that fail to notify or provide feedback to users about how their performance differs from their **[*1019]** appearance potentially deceive users in privacy-implicating ways. Most centrally, this Subsection addresses feedback in general: how to give feedback to users, and enable them to give feedback to _**robot**_ companies in return.

In general, a _**robot**_ should legibly show whether it is protecting or surveilling particular areas. Privacy-sensitive _**robots**_ should be transparent about both what they can sense and how they share information. Some forms of

---

[184] Nissenbaum, supra note 71.

[185] Rueben & Smart, supra note 23, at 33.

notice might be simple. A set of privacy labels on a **_robot_** could be used to give nearby users and third parties some information about what they can detect. These labels could also be broadcast to mobile devices or other screens in an area, giving notice before users are within range of a **_robot_**'s sensors. Labels could even be interactive, allowing users to disable certain sensors or certain forms of sharing. Some research into other forms of **_robot_** transparency has been done, but more is needed. [186]

Transparency alone might not be enough - especially, as discussed, if it is one-time transparency rather than dynamic and ongoing. **_Robots_** sense the world in fundamentally different ways than humans, and it will be hard to always articulate their decisions to human users because sensors are complex technological artefacts. One key challenge is the mismatch between the sensors on a **_robot_** and the five human senses discussed above. A human might specify her privacy preferences as "don't look through the bathroom door when I'm in there." Although this would work for a human observer, a **_robot_** might have a sensor that looks through walls. This mismatch between expectations and the reality of what the **_robot_** can sense is a potential cause for concern, and challenging to address. Other senses beside vision (for example, hearing or touch) need to be explored for privacy protection strategies.

Significant notice and feedback problems arise, as discussed, with anthropomorphic **_robots_**. Additional research is much needed in the field of human-**_robot_** interaction to determine the extent to which **_robots_** can socially engineer increased disclosure by coaxing human users into trusting them, or asking questions in ways that encourage over-disclosure. This will range from studying **_robots_** that invoke more trust than they really deserve, to studying **_robots_** programmed to use interrogation techniques. Going along with our principle of honest anthropomorphism, these kinds of social engineering raise significant privacy concerns.

In closing, although much technology exists that seems viable for creating the first privacy-sensitive **_robots_**, much work still needs to be done in creating whole systems that work in the field and in evaluating the performance of those systems. All the pieces need to be present - not just privacy filters, but the interfaces through which users specify their preferences and **[*1020]** the software framework that decides what to do when filters malfunction. Whole systems like this should be implemented on real **_robots_** and deployed in real-world settings for careful evaluation.

III. A Role for the **_Law_**?

Although significant technical challenges remain, technologies can do important work in mitigating **_robot_** privacy problems. The classic problem for the **_law_** of new technologies is how to embed or enforce important values without creating outdated rules or too strongly disincentivizing or constraining innovation. While not intended to be exhaustive, in this Part we survey what kinds of legal tools might accomplish these goals. We explore whether there is room for the U.S. legal system to encourage, nudge, or even mandate technologists to research and implement these kinds of tools. This is a smaller instantiation of the more general challenges of the relationship between **_law_** and privacy by design.

Using **_law_** to regulate design, technological architecture, or code raises a number of concerns. First, there is the overarching concern that fine-grained regulation of design can constrain technological development. If we tell technologists exactly what to build, they won't explore more creative options, and we will miss out on innovative solutions lawmakers could not foresee. A second concern about using **_law_** to regulate code is with over-enforcement or creating immovable constraints. By building values into design, we may constrain the user experience or too-perfectly enforce the **_law_**, where some space for legal play may be better policy. Third, constraining design through **_law_** can hide the work that **_law_** is doing. A user may fail to realize that her problems with a technology are actually the result of legal policy. Thus, governing through code can be undemocratic. [187]

---

[186] See, e.g., Leah Perlmutter et al., Situated Language Understanding with Human-Like and Visualization-Based Transparency, 12 Online Proceedings of the Conf. on Robotics: Science and Systems (2016), *www.roboticsproceedings.org/rss12/index.html*.

[187] Lessig, Code: Version 2.0, supra note 113, at 133 ("When a government uses other structures of constraint to effect a constraint it could impose directly, it muddies the responsibility for that constraint and so undermines political accountability.").

Governing through code or design also potentially implicates free speech, as indicated in the existing reluctance to impose liability for failed code. United States lawmakers and courts have historically been reluctant to impose liability for failed code when that failure does not cause physical harm. Tort lawsuits against software producers have been successful when the software causes physical harm, but economic loss doctrine precludes a tort claim if there is only economic harm. Similarly, consider the Communications Decency Act ("CDA"). [188] Section 230 of the CDA effectively immunizes online service providers from liability for speech occurring on their platforms. [189] The First Amendment has been found to protect a wide variety **[*1021]** of speech online, potentially including both speech produced by code, and code itself. [190] We do not resolve First Amendment questions here, but consider them as a backdrop for our discussion of potential regulatory tactics. *Robot* design admittedly looks more like product design, and because that design is embodied in a physical object, it may trigger lesser forms of both First Amendment scrutiny and coverage.

What is a lawmaker to do? Creating specific legal rules around *robot* design - whether statutory or regulatory - is a bad idea. Ossifying requirements that can be changed only through legislative process or even regulatory process may create problems in a fast-changing area. Codifying specific design rules may constrain or channel technological development, and could result in technologists missing potentially effective solutions in attempts to obey specific *laws*.

Instead, we suggest an approach of using codified standards that delegate authority to nimbler decisionmakers, with an emphasis on both iteration and process. [191] Legislators could establish general design standards instead of specific rules, delegating interpretation to either courts or agencies that would address specific cases. Or agencies and courts could in turn establish standards or principles that a number of different kinds of design could satisfy. The concern with this approach is that it does not provide much ex ante notice or guidance to technologists. If a *robot* must protect a "reasonable expectation of privacy," how specifically might that requirement be built into its design?

One suggested solution would be to address *robot* privacy at the Federal Trade Commission. [192] The FTC employs its Section 5 authority to regulate unfair and deceptive trade practices, including both data privacy and, more recently, data security. [193] Section 5 settlements already address design to some extent, penalizing deceptive or unfair design choices. In fact, FTC settlement "jurisprudence" has developed a concept of notice that focuses on design rather than on verbal disclosure. [194] Another tool used in Section 5 settlements is requiring companies to employ Comprehensive Privacy Programs (CPPs), which build privacy into a company's processes, including **[*1022]** design. CPPs require iteration and testing, which might be the appropriate approach for enforcing user-centric design. [195]

---

[188] Pub. L. No. 114-38, *129 Stat. 86 (1998)* (codified as amended at *47 U.S.C. § 230* (2012)).

[189] *47 U.S.C. § 230.*

[190] See, e.g., *Zhang v. Baidu.com, Inc., 10 F. Supp. 3d 433, 434 (S.D.N.Y. 2014)* (characterizing search engine algorithmically produced results as "in essence editorial judgments about which political ideas to promote"); see also *Langdon v. Google, Inc., 474 F. Supp. 2d 622 (D. Del. 2007);* see also Robert Post, Encryption Source Code and the First Amendment, *15 Berkeley Tech. L.J. 713 (2000).*

[191] See, e.g., Mulligan & King, supra note 6, at 1034; Rubinstein, supra note 6, at 1453.

[192] Hartzog, supra note 5, at 788.

[193] Solove & Hartzog, supra note 101, at 598.

[194] Hartzog, supra note 5, at 816 (discussing the "4Ps": prominence, presentation, placement, and proximity); Mulligan & King, supra note 6, at 1029.

[195] Mulligan & King, supra note 6, at 1028.

The oft-voiced concern with this approach is that in the early stages of **_robot_** development, enforcing broad standards of unfairness or deception might not give adequate notice to technologists of what behavior is illegal. [196] To address this concern, as it has with other new technologies, the FTC could issue guidance on **_robot_** design (perhaps incorporating our five principles) and host workshops. [197] (Some also contend that the FTC approach does in fact provide significant notice to companies, with FTC settlement agreements functioning as a type of common **_law_**. [198]) While FTC guidance would not itself be enforceable, it could nudge industry standards in a particular direction. Once industry standards are established, this could in turn provide the basis of future FTC enforcement actions against **_robot_** designers who fail to implement standard privacy protections. State attorneys general could enforce similar state **_laws_** in the name of consumer protection, to prevent particularly egregious privacy practices. [199]

In addition to settlements, guidance, workshops, and nudging of industry best practices, the FTC could use some tools it has not yet employed. The FTC could start requiring companies to maintain design documents. It could establish performance standards, such as for user comprehension of **_robot_** disclosures or around user susceptibility to anthropomorphic manipulation. One scholar has argued that at least in the area of children's privacy, the FTC has statutory authority to promulgate such standards as a rule. [200] (Though, as discussed above, it is unclear whether **_robots_** or how many **_robots_** would be subject to the COPPA regulatory regime.) Companies could be monitored for compliance and update their performance standards as testing results come in. [201]

Another possible venue for addressing **_robot_** privacy-by-design is the National Telecommunications and Information Administration ("NTIA") at **[*1023]** the Department of Commerce. [202] The NTIA has hosted, in recent years, a series of multi-stakeholder meetings aimed at producing best practices on notice, privacy, or security in a variety of subject matter areas: mobile applications, facial recognition, drones, and the Internet of Things. [203] One of us has criticized the efficacy of this process, asserting that in the absence of federal data privacy **_law_** or other penalties, private actors are reluctant to meaningfully contribute to substantive privacy best practices. [204]

Moving beyond federal agencies, courts might interpret existing state privacy **_laws_** as applicable to home **_robots_**. For example, a court might find that a **_robot_** that looks through walls has in fact committed the intrusion tort. Case-by-case assessment of liability under existing legal regimes has the benefit of moving incrementally, and being fact-specific. It also has the benefit of potentially functioning like a performance standard: those technologies that fail to deliver adequate privacy protections would lose in individual cases, sending technologists back to the drawing board. However, there are multiple limitations to existing privacy torts, especially in their application to data privacy. Courts often (but not always) find that sharing information with one person eliminates a privacy interest. Users who employ **_robots_** that share information with third parties outside the home may face similar hurdles.

---

[196] See Geoffrey Mann, Time for Congress to Cancel the FTC's Section 5 Antitrust Blank Check, Tech. Liberation Front (Dec. 20, 2012), _https://techliberation.com/2012/12/20/time-for-congress-to-cancel-the-ftcs-section-5-antitrust-blank-check/_; Berin Szoka, To Protect Consumers, Congress Must Limit the FTC's Discretion, The Hill (Aug. 27, 2015, 2:00 PM),_http://thehill.com/blogs/congress-blog/251934-to-protect-consumers-congress-must-limit-the-ftcs-discretion_.

[197] Hartzog, supra note 5, at 827-28; Mulligan & King, supra note 6, at 1030; Rubinstein, supra note 6, at 1447.

[198] See generally Solove & Hartzog, supra note 101.

[199] Citron, supra note 104, at 749.

[200] See Lauren E. Willis, Performance-Based Consumer **_Law_**, _82 U. Chicago L. Rev. 1309, 1375 (2015)._

[201] See _id. at 1373-76._

[202] See generally Kaminski, supra note 112, at 925.

[203] Id.

[204] Id.

28

Of course, Congress or state legislators could craft **_robot_**-specific **_laws_**. If this takes place at the state level, then some experimentation with different regulatory regimes would be allowed. [205] The risk with technology-specific regulation, like the risk with writing specific design rules, is that it may become outdated. [206] Or it might fail to address problems with data privacy as a whole while focusing on just one technology. The United States already takes a piecemeal approach to privacy protection, although it tends to focus on specific sectors at the federal level rather than specific technologies. Adding new technology-specific **_laws_** to the existing tangle would add to regulatory costs, which could preclude smaller market entrants.

If either state legislatures or the U.S. Congress were to pass a **_robot_**-specific privacy **_law_**, there are significant questions about what it might look like. It would be dangerous to enshrine particular technical requirements into a **_law_**, given the pace of technological development and the fact that better solutions may be arrived at in the future. Lawmakers could instead adopt a standards approach, putting in place more general requirements or principles, **[*1024]** like the five we outline here. Courts or agencies could then interpret these principles, or could work with private actors to create subject matter-specific best practices. This two-system approach - enforcing the broad principles, or giving private actors the option of clarifying the **_law_** by articulating how those principles might be specifically applied - was the concept advanced by the proposed Consumer Privacy Bill of Rights. [207] As long as there is a significant enough threat of enforcement, this approach might spur companies to arrive at design policies that both protect users and do not unduly constrain the technology.

IV. Conclusion

Robots in the home pose a variety of privacy threats. We identify three broad types: data privacy harms, boundary-management harms, and harms caused by the way **_robots_** are designed to socially interact with humans. Current U.S. **_law_** is not well-equipped to address these problems. Accordingly, we encourage both technologists and regulators to approach these concerns by considering **_robot_** design. We propose five principles for designing privacy-sensitive **_robots_**: data minimization, purpose specification, use limitation, honest anthropomorphism, and dynamic participation and feedback. These principles are derived from the FIPPs, but we note that the FIPPs alone do not adequately address all the privacy harms we identify here.

We hope to encourage technologists to both adopt particular existing privacy-protective technologies, and research and develop further technologies to mitigate these harms. Similarly, we hope to enable potential regulators to understand both the toolkits available to technologists and the limitations in what the technology can actually solve. Finally, we encourage regulators to choose more dynamic, nimbler forms of regulation instead of requiring particular technological specifications or writing technology-specific **_law_** at this stage. Bill Gates may have been correct that there will eventually be a **_robot_** in every home. [208] But unless they learn to avert their eyes, home **_robots_** may not be very welcomed or trusted by their owners.

End of Document

---

[205] Margot E. Kaminski, Drone Federalism: Civilian Drones and the Things They Carry, 4 Calif. L. Rev. Cir. 57 (2013); Paul M. Schwartz, Preemption and Privacy, _118 Yale L.J. 902, 916 (2009)._

[206] See generally Michael Birnhack, Reverse Engineering Informational Privacy **_Law_**, _15 Yale J.L. & Tech. 24 (2012)._

[207] Dennis D. Hirsch, Going Dutch? Collaborative Dutch Privacy Regulation and the Lessons It Holds for U.S. Privacy **_Law_**, 2013 Mich. State L. Rev. 83, 86-87; Kaminski, supra note 112, at 926.

[208] See Gates, supra note 8.