

Themes and Research Directions in Privacy-Sensitive Robotics

Matthew Rueben¹, Alexander Mois Aroyo², Christoph Lutz³, Johannes Schmölz⁴,
Pieter Van Cleynenbreugel⁵, Andrea Corti⁶, Siddharth Agrawal⁷, and William D. Smart⁸

Abstract—Privacy is crucial for healthy relationships, but robots will impact our privacy in new ways—this warrants a new area of research. This paper presents work from the first workshop on privacy-sensitive robotics. We identify the seven research themes that should comprise privacy-sensitive robotics research in the near future: data privacy; manipulation and deception; trust; blame and transparency; legal issues; domains with special privacy concerns; and privacy theory. We intend for the research directions proposed for each of these themes to serve as a roadmap for privacy-sensitive robotics research.

I. INTRODUCTION

All humans need privacy, although different cultures seek it in different ways [1]. Privacy is a crucial requirement for relationships to flourish [2] and for individuals to mature and have freedom. Robots are capable of violating human privacy: they can collect and share information, move through personal spaces and territories, and interact with people socially [3]. Thus, robots and the companies that design them will need to earn people’s trust if they are going to be accepted.

The field of human-robot interaction (HRI) has barely begun to think about privacy. Privacy scholars have formulated several theories of how privacy works between humans [4][5], but we do not know if they extend to interactions with robots. Plus, privacy has several distinct dimensions [6] and is often the subject of legal debates—privacy is a difficult and complex issue. HRI researchers have neither drawn significantly from privacy scholarship nor designed practical systems to comprehensively protect user privacy even though techniques exist to do so, e.g., from computer vision to filter out parts of videos or from autonomous navigation to avoid private rooms. We propose to draw on privacy scholarship and push privacy studies forward in HRI by founding a new research area called “privacy-sensitive robotics.”

¹Matthew Rueben is a Ph.D. Candidate in the Robotics Program at Oregon State University

²Alexander Mois Aroyo is a Ph.D. Student in the Department of Robotics, Brain and Cognitive Sciences at Istituto Italiano di Tecnologia

³Christoph Lutz is Assistant Professor in the Department of Communication and Culture at the BI Norwegian Business School, Oslo

⁴Johannes Schmölz is a Ph.D. Student in Electrical Engineering and Computer Science at Coburg University

⁵Pieter Van Cleynenbreugel is with the Faculty of Law, Political Science, and Criminology at the University of Liège

⁶Andrea Corti is with Technische Bürgernetzwerk Austria (tbn^a)

⁷Siddharth Agrawal is an undergraduate student in Mechanical Engineering at the Indian Institute of Technology, Delhi

⁸William D. Smart is Professor in the Collaborative Robotics and Intelligent Systems (CoRIS) Institute at Oregon State University

To get more researchers involved in this new and growing area, a workshop was held at the ACM/IEEE International Conference on Human Robot Interaction (HRI) on March 6th, 2017 in Vienna, Austria. The workshop, entitled “Privacy-Sensitive Robotics,” brought together researchers from a wide variety of communities to help define privacy-sensitive robotics by identifying the important research questions. The 15 attendees included roboticists, social scientists, and legal scholars, all with very different backgrounds. This paper is a direct product of that workshop as well as additional work that followed to organize and develop those thoughts.

This paper will define the pressing issues in privacy-sensitive robotics from a variety of perspectives, and will suggest a research agenda to address these problems. We begin by introducing the concept of privacy as understood by scholars (Section II) and listing published work that we would consider to be privacy-sensitive robotics research (Section III). We then identify 7 research themes within privacy-sensitive robotics (Sections V–XI): (1) data storage, processing, and filtering, (2) how robots can trick people into giving up personal information, (3) trust, (4) blame, (5) privacy regulations and other legal topics, (6) special private domains like the home, and (7) privacy scholarship outside of HRI. For each theme there is a brief discussion followed by concrete suggestions for future research.

II. BACKGROUND ON PRIVACY SCHOLARSHIP

Privacy scholars have argued about the definition and scope of “privacy” since the late 19th century [7]. Moore defines privacy as, “control over access to oneself and information about oneself” [8]. This is a “control-based” definition of privacy, in which it doesn’t matter whether somebody accesses you or your information, but rather whether you can control that access. Altman’s theory defines privacy as a *boundary regulation process* wherein people try to achieve their ideal privacy state by using certain *mechanisms* to regulate interaction with others [9]. Notice how this definition allows privacy to sometimes mean *more* interaction with others, and sometimes *less* interaction. Austin [10] defines privacy as freedom from “public gaze,” which includes video surveillance in public. Nissenbaum’s approach to privacy, “contextual integrity,” focuses on the idea that different norms of information gathering and dissemination are observed in different contexts [5].

Another way to define privacy is through a taxonomy that covers all the different parts of privacy. Rueben et al. have compiled one from the privacy literature [6] whereas Koops

et al. [11] have made a typology based on constitutional privacy protection in nine different countries. The taxonomy by Rueben et al. [6] found that privacy literature divides privacy into (1) Informational privacy, over personal information; (2) Physical privacy, over personal space or territory; (3) Psychological privacy, over thoughts and values; and (4) Social privacy, over interactions with others and influence from them.

III. RELATED WORK IN PRIVACY-SENSITIVE ROBOTICS

A few studies have been published that we would consider privacy-sensitive robotics research. Syrdal et al. [12] studied disclosure of personal information and Caine et al. [13] studied privacy-enhancing behaviors. Wong & Mulligan [14] studied what concept videos can communicate to potential users about privacy. Denning et al. [15] demonstrated security vulnerabilities in commercially available robots. Lee et al. [16] studied users' perceptions of privacy with a social robot in the workplace. Hubers et al. [17], Butler et al. [18], and Klow et al. [19] studied the trade-off between filtering the robot's video feed to protect user privacy and keeping it unfiltered so the remote operator can use the robot effectively. Rueben et al. [20] studied whether physical, persistent interfaces provide benefits over GUIs for specifying user privacy preferences, and Wagner [21] presents an architecture for automatically detecting private objects and locations. Rueben et al. [22] studied the effects of interpretive frames, Tonkin et al. [23] of embodiment, and Vitale et al. [24] of transparency on privacy judgments. Also, Krupp et al. [25] studied what privacy concerns people have about telepresence robots.

There has also been some theoretical work in privacy-sensitive robotics. Legal scholar Calo's chapter on privacy in *Robot Ethics* [3] identifies three ways that robots present new privacy concerns: direct surveillance, increased access, and social meaning. Another paper by Calo focuses on how drones [26] might prompt changes in U.S. privacy law, and legal scholar Kaminski comes to a similar conclusion in her discussion of home robots and privacy [27]. Kaminski et al. [28] present several categories of potential privacy harms by robots, then some technological as well as legal solutions. Lutz & Tamò call for a new class of jobs to bridge the divide between privacy regulators and engineers [29], and have also argued for the usefulness of actor network theory (ANT) in their analysis of privacy in healthcare robotics [30]. Shulz & Herstad [31] apply the privacy framework by Palen and Dourish to a mobile robot in the home; similarly, Sedenburg, Chuang, & Mulligan [32] apply the Fair Information Practice Principles (FIPPs) as well as research ethics to the development of therapeutic robots in the home.

Rueben & Smart [33] have also presented a roadmap for privacy-sensitive robotics as a new field, albeit from an HRI perspective and without any claims of completeness. In this paper we tried to cover *all* the significant themes within privacy-sensitive robotics by including researchers from several different disciplines. Our goal was to generate

a wide variety of research directions that could be followed by many researchers in parallel.

IV. THE FIRST WORKSHOP ON PRIVACY-SENSITIVE ROBOTICS: PROGRAM AND GOALS

The workshop was held in conjunction with the ACM/IEEE International Conference on Human-Robot Interaction (HRI) on March 6th, 2017 in Vienna, Austria. It was organized by Matthew Rueben, Bill Smart, Cindy Grimm, and Maya Cakmak. The goal of the workshop was to bring together the people who are starting to do research on privacy in human-robot interaction and begin to form a community. The deliverable was to map out all the important research areas in privacy-sensitive robotics, identify pressing research questions in each area, and compile these into a roadmap to help focus the efforts of this new community.

The workshop comprised two invited speakers, a set of shorter contributed talks, and an extensive ideation session wherein workshop participants brainstormed productive research directions, both from their own prior work and in response to the work presented at the workshop. The invited speakers were Woodrow Hartzog, a legal scholar now at Northeastern University, and Christoph Lutz, a communication and culture scholar at the BI Norwegian Business School. The 15 participants included human-robot interaction researchers with a variety of backgrounds as well as people from privacy & security, human-computer interaction, and law. A full schedule is available along with the contributed papers on the workshop website¹. The 7 research themes presented in this document (Sections V–XI) and the proposed research directions therein are the result of this collaborative and interdisciplinary process.

V. THEME 1 OF 7: DATA PRIVACY

A. Storage and Processing

Personal information may be collected, processed, stored and shared by robots and the people who have access to their hard drives. Robots can be divided into three categories depending on how personal information is processed and stored:

- *Onboard processing*: robots such as the Roomba are able to do all their information processing and storage within their body, without the need of external components². Onboard processing seems to be the best solution for privacy, but it may offer low performance due to its technological limits.
- *Local processing*: robots such as Cozmo need a local computing resource like a smartphone or PC to function, although no Internet connection is required³.
- *External processing*: robots such as Pepper or Hello-Barbie need an external resource to function⁴. Such

¹<https://sites.google.com/oregonstate.edu/hri-2017-privacy-workshop/program>

²<https://www.technologyreview.com/s/541326/the-roomba-now-sees-and-maps-a-home/>

³<https://support.anki.com/hc/en-us/articles/236021007-COZMO-Cozmo-Basics>

⁴<https://toytalk.com/hellobarbie/terms/>

resources are usually located in the cloud and used by the company to process the data. Transmitting the data via the Internet exposes it to additional security risks. Also, cloud servers could be located anywhere in the world, and might be owned by a third party company that provides cloud storage or processing as a service. Users should be notified about which data is transmitted and stored externally and the possible risks of doing so.

The growing field of cloud robotics raises some additional privacy concerns. One application of cloud robotics is in remote robot learning [34], in which a robot is controlled by a remotely located user to teach the robot to perform tasks that are difficult to do autonomously, such as grasping. Also, some of the robot supervising roles can be outsourced to places where more human capital is available. Both of these aspects of cloud robotics pose increased privacy and security risks. If unauthorized people were able to get control of the robot, they could cause physical damage or spy on local users. Willow Garage’s Heaphy project on remote robot learning ended partly due to these privacy and security challenges⁵. Privacy and security researchers are needed to help analyze all the options discussed in this section and especially to identify any special risks from factors that are unique to robots, such as embodiment and mobility.

B. Technical Enforcement

One way to enforce privacy protection is to create levels of clearance and to allow the robot to recognize and categorize people into these levels. For example, if the robot lives with a family all the family members could have permission to access family-related information. If a guest comes to the house, the robot should recognize him or her as an outsider and should not disclose that kind of information. If new information is added to the robot, it should be stored at the strictest clearance level by default.

Some privacy concerns can be lessened or even avoided completely via technical solutions. For example, the robot could avoid certain areas, perhaps during specified hours of the day, in case it sees something private. Perhaps the bedroom would be off-limits at night. A second type of technical solution relies on detection or recognition: certain objects, classes of objects, or people might be designated as needing to be blurred out whenever they are detected. For example, maybe a concern about the robot seeing secret documents could be addressed by blurring out anything detected as text. For security the blurring should happen at the lowest level possible, such as in the firmware of the camera where it cannot be accessed by the robot’s operating system. Finally, a third class of technical solutions would take a “privacy by design” approach by mounting the robot’s camera too low to see the tops of tables and desks where documents tend to be. There is a tradeoff here, though, since it might need to see things up high to do certain tasks. For future work we recommend to implement each of these

solutions on a real robot and do exploratory testing to identify broad problems.

C. Data Preferences

In many situations, users will be asked to tell the robot their privacy preferences: which data to collect, where to go, and what to filter. Users might want to adjust these preferences on the fly as issues arise or situations change, perhaps even modifying or deleting data that has already been collected. In future work, interfaces for specifying visual or spatial privacy preferences should be designed. Such interfaces could be like those evaluated by Rueben et al. [20], or they could use other modalities like spoken dialogue, gestures, or drawing. It will be important for these interfaces to accommodate subtle changes to privacy preferences over time.

A problem arises about how to classify information so the robot knows which information can be shared and with whom. For example, robots might talk to humans or amongst themselves, perhaps to gain more information for learning, and might say things that another user would have liked to keep private. A possible solution could be to use an “opt-in” regime, so the robot won’t collect or share certain personal information unless you explicitly give your permission (“opt in” to it). Another approach is for the robot to recognize different contexts automatically and respect the different privacy rules for each context as well as for moving between contexts—i.e., to use Nissenbaum’s idea of contextual integrity [5]. Building this sort of Nissenbaumian framework would be completely new research for human-robot interaction.

D. Personalization, Learning, and Inference

Social robots might observe our behavior and adapt to it over time. This behavioral data might even be collected in a private (e.g., home) setting. Perhaps a family’s daily routine would be observed by the robot and examined for patterns using machine learning techniques. Since the interpretations of the private information collected by the robot are not only stored but embodied—e.g., if the robot follows along with the daily routine—a violation of the user’s privacy could happen. Traits and characteristics of family members or even of the whole family might be visible in the behavior of the robot in front of visitors or strangers. Maybe the robot does something impolite during a dinner party that it had clearly learned from a family member. Future work could enable robots to understand which behaviors that were learned in one context are inappropriate to display in other contexts because of the inferences people could make. Alternatively, the robot could simply revert to its default, unpersonalized behaviors around strangers.

VI. THEME 2 OF 7: MANIPULATION AND DECEPTION

Humans use different behaviors and personae depending on whether they are with family, friends, coworkers, or strangers. Like humans a robot should be able to adapt its personality to different people and situations to improve

⁵<https://spectrum.ieee.org/automaton/robotics/robotics-software/the-heaphy-project>

its relationships. For example, a social robot may act more familiarly with a friend than with a stranger. But this also opens the door for robots to be manipulative or deceptive by pretending to be something they are not. This is a privacy risk inasmuch as it could give the robots supervisors access to users' personal information. The question is: how is this process different (if at all) for manipulative robots than for manipulative humans?

A. Social Engineering (SE) using Robots

It is possible that robots could be used to trick, con, or dupe humans using social engineering (SE) techniques [35]. For example, Booth et al. [36] showed that robots could use social engineering techniques to sneak into a campus building. Other kinds of attacks we foresee include hijacking robots and using them to surveil an area that can be attacked or robbed in the future, or just disabling the robot's cameras so a crime can occur without being recorded. Research on social robots is still very young, so lots of work is still needed to understand which social engineering techniques could be performed by robots, perhaps including new ones that humans cannot perform.

B. Security Vulnerabilities

Fast-paced development of complex systems can leave behind security holes. If a malicious person hacks into a robot there are particularly severe privacy risks because of both the sensors onboard and the robot's ability to move, perhaps in a private setting. There have been some recent examples of security vulnerabilities in robots being exploited to violate privacy. The Hello Barbie doll was hijacked [37] and turned into a surveillance device to spy on children—it recorded the conversations and sends them via WiFi to the internet. Similarly, teddy bears could be used to spy in houses or talk to children [38]. Research in this area might include designing architectures that are especially careful to protect privacy-relevant features of the robot such as sensor feeds, stored data, and motor control. Researchers should also look into limiting the data that gets stored or transmitted, or transmitting only the information that is necessary instead of full-fidelity, raw data [39].

C. Education

One way to mitigate the risk of maleficent deception would be to promote education and experience with robots over time. Robots are still relatively new to our society; when computers were still new, old types of scams were quite successful, like the Nigerian Prince advance fee scam. Now, through education and media, the success rate of those types of scams is relatively low—perhaps education and awareness-raising can do the same for scams that use robots. Even the simple lesson that robots are built around computers and therefore could have security vulnerabilities could prevent a lot of harm.

People should also know their rights. In the EU, all robots that collect personal data will fall under the General Data Protection Regulation (GDPR), which mandates that users

be informed of how their data will be used before they consent to data collection. In the US, a deceptive robot could be regulated by the Federal Trade Commission (FTC) with similar requirements [40]. Developing educational programs about privacy law and risk management with robots is an almost untouched research area.

VII. THEME 3 OF 7: TRUST

The topic of trust is closely related to privacy concerns. Richards & Hartzog [41] have argued that we should not just focus on privacy harms, but also on how privacy can enable trust. Rousseau, Sitkin, Burt, and Camerer [42] have defined trust as “a psychological state comprising the intention to accept vulnerability based upon positive expectations of the intentions or behaviors of another” (p. 395). Trust has several sub-dimensions, including one's general disposition towards trusting people (e.g., strangers), trust of institutions, trust towards individual people based on what one knows about them, and the willingness to engage in a transaction based on that trust [43]. With regards to robots, research has started to investigate user trust through qualitative, quantitative and experimental methods. Hancock et al. [44] review much of this research in their meta-analysis of which factors affect trust in human-robot interaction.

Future research should continue looking for links between privacy concerns and trust in human-robot interaction because trust is so important for a successful interaction. In addition, researchers should study the factors that affect users' trust that a privacy-protecting system is actually protecting their privacy. Whether it's trust of the software, the manufacturer, or the robot itself, the privacy protection system is not nearly as useful if people aren't confident that their privacy is actually being protected. Finally, future work should study how privacy concerns intersect with trusting relationships like between a patient and a (robot) doctor or if the robot is treated as a family member and is expected not to gossip about what happens at home.

VIII. THEME 4 OF 7: BLAME AND TRANSPARENCY

One of the core questions in privacy-sensitive robotics is whom to blame when privacy violations occur. This is especially difficult when the robot's behavior is at least partially influenced by a remote operator or supervisor, or the robot is capable of learning, or has behavior that is difficult to predict for some other reason. Here we consider several parties that could be blamed: designers/manufacturers, owners, distributors/controllers, and the robot itself.

- 1) *Designers or manufacturers* may be blamed for having provided improper or faulty privacy protections on their (learning) robot. Alternatively, since the manufacturers cannot control the stimuli from which the robot learns, one could argue that a robot designed, produced or programmed in accordance with privacy standards would not make the manufacturer liable for privacy breaches resulting from a robot's learning capacities. Future work should study existing product liability law regimes across different jurisdictions and whether they

need to be changed to better handle robots that can learn.

- 2) *Owners* may be liable for mistakes made by their robots. Robots with learning capabilities can to some extent be compared to children or pets, for which owners often are liable even when the owner himself committed no fault [45]. Could users be taught that they also have a responsibility for “raising” the robot, and could manufacturers use disclaimers to escape liability in this regard? If so, when (if ever) has the robot “grown up”, and who is responsible then?
- 3) Data assembled by a robot may be transferred to a certain *controlling company*, especially whenever owners have consented to this transfer in general terms and conditions. That company may sell your personal data and allow other businesses to use the robot to give you personalized suggestions. As the robot collects that data it is important to clarify to what extent data controlling companies must respect the privacy of certain information and whether robots need to be programmed to only transfer certain kinds of information.
- 4) Questions remain when and to what extent *the robot itself* is to be considered a moral agent that is liable for its actions. At some point, the robot may have reached the maximum of its learning abilities and the question can then be asked whether owners are still liable for privacy-breaching actions at that point.

It is also important to be *transparent* about who is actually watching you and recording or using your data, especially since the answer will probably affect your level of privacy concern. Are people more (or less) conscious of their privacy if they know who is behind the robot interpreting the data? When a robot is autonomous, should information about its programming be given? More ethical and sociological research is needed here, setting up studies about how people make privacy-related attributions in instances of shared control over robots. The results of those studies can help us develop features that increase transparency about robotic systems that handle personal information.

IX. THEME 5 OF 7: LEGAL

A. Robots as Persons or Family Members

As social robots enter the home they could be treated like members of the family in a similar way to how pets are sometimes treated. Over time, the humans might build up trust towards the robot so that they speak and act freely around it, expecting that it will respect their privacy at home. This trust is a privacy risk if robots are subject to search by law enforcement with a proper warrant, access by law enforcement without a warrant under the third party doctrine⁶, or perhaps robots will even be mandatory reporters of any signs of domestic abuse. This is especially problematic

⁶In the US, one can have no reasonable expectation of privacy for information that has been given to third parties such as banks or internet service providers ... or, perhaps, robots.

because robots often have the sensing capabilities (e.g., cameras and microphones, plus the ability to move around the house) to assemble a very detailed picture of home life. One solution would be to consider the robot a close family member with evidentiary privilege in court—i.e., the robot would not have to testify against the family. This might make sense as a way to avoid chilling effects on expression and behavior inside the home—Australia grants evidentiary privilege to the parent-child relationship for similar reasons [46]. A more extreme solution would be to consider robots as “electronic persons” as proposed by the European Parliament⁷ instead of as objects of search, forcing police to go through some sort of interrogation process instead of having access to *everything* stored on the robot. Research by legal scholars is needed to work out the details for each of these options and to consider which would be the most reasonable.

B. Regulating Robotics Companies

Should legal obligations be imposed on businesses to include privacy-sensitive features in their robotic products? Is it legal to compel manufacturers to build in privacy features limiting recording and storage capacities of robots to protect the privacy of their owners? Privacy protection does not always align with business interests because the features that customers want for their robots often require collecting, processing and storing personal information. In the European Union, the right to data protection and especially the right to be forgotten may impact the way robots are designed—more research is needed to determine how.

An alternative to government regulation would be to incentivize businesses to self-regulate. One possible solution would be through a technical standard created by the International Standards Organization (ISO). Standards are adopted without government interference and at the same time bring confidence that a product is safe, reliable and of sufficient quality. Research in this area could determine the feasibility and usefulness of creating a standard for privacy-sensitive robots.

C. Privacy Education for Users

Should there be a legal obligation to provide privacy education to robot users? If so, who should be responsible for providing it? Public education plans could mandate teachings on what to say and what not to say to robots, and could be applied to children from a young age. Solove & Hartzog have argued that the US Federal Trade Commission (FTC) and its equivalents in other countries may already have the powers to intervene in this field [47]. Research questions arise about what ways consumer protection agencies can play a role, whether they have sufficient intervention powers or whether their actions should be accompanied by a more developed legal framework. The European Union has proposed the establishment of a robotics agency playing this role⁸.

⁷See European Parliament, resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL))

⁸European Parliament, Resolution of 16 February 2017, para 15 and onwards.

X. THEME 6 OF 7: PRIVATE DOMAINS

Some domains have special privacy concerns that warrant their own, additional research efforts. Here we survey three such domains and the special challenges of deploying robots in each.

A. *Robotic Surgeons*

Use of robotic surgeons such as the daVinci surgical system allows surgeons to perform very precise operations inside the body [48]. These systems have already been used to conduct more than a million operations worldwide. Although the robotic surgeons being used right now are teleoperated, advancements in computer vision and robotics have made it possible for the robots to carry out some repetitive surgical tasks autonomously. One way to carry out such delicate tasks would involve collecting large sets of sensitive, personal data, including images of a person's naked body. It is essential to have very clear imagery of the environment during surgery, so it becomes difficult to use privacy filters like those proposed by Hubers et al. [17] and Butler et al. [18]. Utility will supercede privacy during surgery since human lives are at risk. It is important, though, to use very secure network connections for communication between the surgeon and the robot, as people could break in and steal this highly sensitive data. It would be even worse if someone were able to get control of the robot and send commands that cause it to harm the patient.

B. *Robotic Nurses and Caretakers*

Robotic nurses and caretakers could provide a cheaper alternative to human caretakers for providing care in hospitals or even in homes. Robots could help people with disabilities, or older adults who want to age in place [48]. To operate, these robots will need to collect private information such as images of the person and of the environment, maps of the house, and medical information. This raises plenty of concerns even outside of healthcare applications, but robotic nurses and caretakers might be treating people who are especially vulnerable to privacy harms: people who might not be strong enough to resist invasive behaviors or mentally sound enough to understand what is happening. Plus, patients might be embarrassed about their appearance or the medical procedures they undergo, so it is more likely that these robots will capture information that should be kept private. Note that many of these same concerns apply to robotic toys designed for children.

C. *Robots in the Home*

Denning et al. [15] conducted a study exposing security vulnerabilities on three household robots: the WowWee Rovio, the Erector Spykee, and the WowWee RoboSapien V2. Domestic robots like the Roomba and Jibo offer great utility for homes, but with their wide range of sensors they also get access to a lot of private information. For example, the Roomba 960 vacuum robot can build a complete map of a home and can interface with Alexa and the Google

Assistant⁹. This integrated network of home devices could be useful, but also more vulnerable to privacy and security threats. What if the companies that manufacture these robots were to sell this private data to other companies that you don't trust? There is also ambiguity about the use of this data; for example, what happens when the previous owner of a house decides to sell their Roomba's map of the house to someone else without the consent of the current owner?

XI. THEME 7 OF 7: THEORIES AND PERCEPTIONS OF PRIVACY

We need to understand privacy if we are going to build privacy-sensitive robots. Several systematic reviews have been published recently about privacy research in the social sciences [49][50][51] and more particularly in communication and information systems [52]. This section is dedicated to the areas of general privacy research (i.e., no robots required) that will be most important and inspirational to privacy-sensitive robotics.

A. *Theories or Models of Human Privacy*

In addition to economic and legal scholarship, privacy as a social norm has been an important topic in sociology and communication. Prominent theories developed in this area include communication privacy management theory [4], privacy as contextual integrity [5], networked privacy [53], and privacy by design [39], although the latter is also connected to more technical disciplines.

These theories can help us build frameworks for privacy protection. For example, a framework inspired by Nissenbaum's [5] contextual integrity would use the idea of appropriateness and distribution rules in different contexts, whereas communication privacy management theory [4] might inspire the idea of privacy boundaries being modeled as "thick" vs. "thin" depending on how bad it would be if that boundary were breached. Robots equipped with theory-inspired privacy protection frameworks could then be used to test those theories via user studies. This sort of work should promote collaborations by HRI researchers with privacy researchers in law, philosophy, the social sciences, privacy and security experts, and computer scientists.

B. *The Subjective Value of Privacy*

How much do individuals value privacy relative to other things such as convenience or safety? Research on the *privacy paradox* has shown that in many scenarios users value privacy but are quick to give it up for short-term monetary or social rewards. Dinev & Hart [54] postulate that users perform a privacy calculus by consciously weighing the benefits of a transaction or service against its privacy risks. Privacy is in that sense like a commodity [55] that can be traded in against a benefit such as access to an affordable ride through Uber or to potential dating partners through Tinder. Despite this, people—including younger people—still value their privacy: survey results have shown little

⁹<http://www.irobot.com/For-the-Home/Vacuuming/Roomba.aspx>

difference between adults and minors in their concern about privacy [53].

How much people value privacy can influence how willing they are to engage with robots that collect personal information. Very little research has studied the subjective value of privacy with regards to robots and their usefulness—a rare example is the study by Butler et al. [18] that first used the phrase “privacy-utility tradeoff” in HRI.

We would recommend future studies in HRI to look at the value of privacy from different methodological standpoints. Ethnographic and observational studies in natural settings could look at how users of social robots trade off privacy for other things like dependability or personalization. Making these studies longitudinal could complement the many studies done on the privacy paradox, which barely look at developmental trajectories over a longer period of time. Also, theoretical approaches such as actor-network theory (ANT) and science and technology studies (STS) could be used to look at how robot engineers and manufacturers build their systems so as to encourage users to prioritize certain values. Similarly to the field experiment by Beresford et al. [56], different “invasiveness” scenarios (a robot collecting more data or accessing more personal rooms) could be combined with different utility levels (a robot offering more or less useful services to an individual) to test the privacy-utility tradeoff.

C. Privacy Taxonomies

Privacy is an umbrella term that refers to several distinct concepts. According to Solove [57], “privacy is a sweeping concept, encompassing (among other things) freedom of thought, control over one’s body, solitude in one’s home, control over personal information, freedom from surveillance, protection of one’s reputation, and protection from searches and interrogations.” (p. 1). Rueben, Grimm, Bernieri, & Smart [6] have provided a taxonomy of sub-concepts of privacy based on a review of privacy literature, whereas Koops et al. [11] have made a typology based on constitutional privacy protection in nine different countries. Privacy-sensitive robotics should be concerned with all the different types of privacy, since modern robots have the capacity to collect information, reach into physical spaces and rooms through their mobility, affect humans’ psychological states through interaction, and create social bonds [3]. Privacy-sensitive robotics researchers should test whether all the sub-concepts of privacy apply to human-robot interactions since the literature was written about privacy between humans. A refined taxonomy would be a useful tool for designing and testing privacy-sensitive robots.

XII. CONCLUSION

In this paper we have presented an overview of open research areas in privacy-sensitive robotics. We clustered these areas into 7 themes: (1) data storage, processing, and filtering, (2) manipulation and deception, (3) trust, (4) blame, (5) privacy regulations and other legal topics, (6) special private domains like the home, and (7) privacy scholarship

outside of HRI. It should be clear from the diversity of these research areas that privacy-sensitive robotics research will require collaboration across many different disciplines. We hope this paper will be used as a roadmap by privacy-sensitive robotics researchers, so we have focused on research directions that could use attention in the near future.

Privacy-sensitive robotics research is young, but growing quickly. We hope to encourage and guide this growth. If robots are to be accepted and useful in our society, they will need to respect our personal privacy. Beyond mere acceptance, however, good privacy policies and technologies will help avoid chilling effects on society, like being afraid to express ourselves freely at home or to go out in public. We must protect the silence, reflection, and private conversations that are such important ingredients for human flourishing.

REFERENCES

- [1] I. Altman, “Privacy Regulation: Culturally Universal or Culturally Specific?” *Journal of Social Issues*, vol. 33, no. 3, pp. 66–84, 1977.
- [2] J. C. Inness, *Privacy, intimacy, and isolation*. Oxford University Press, 1992.
- [3] R. Calo, “Robots and privacy,” *Robot Ethics: The Ethical and Social Implications of Robotics*, Patrick Lin, George Bekey, and Keith Abney, eds., Cambridge: MIT Press, 2010.
- [4] S. Petronio, “Communication boundary management: A theoretical model of managing disclosure of private information between marital couples,” *Communication Theory*, vol. 1, no. 4, pp. 311–335, 1991.
- [5] H. Nissenbaum, “Privacy as contextual integrity,” *Wash. L. Rev.*, vol. 79, p. 119, 2004.
- [6] M. Rueben, C. M. Grimm, F. J. Bernieri, and W. D. Smart, “A taxonomy of privacy constructs for privacy-sensitive robotics,” arXiv:1701.00841v1 [cs.CY].
- [7] J. DeCew, “Privacy,” in *The Stanford Encyclopedia of Philosophy*, Fall 2013 ed., E. N. Zalta, Ed., 2013.
- [8] A. D. Moore, “Privacy: its meaning and value,” *American Philosophical Quarterly*, pp. 215–227, 2003.
- [9] I. Altman, *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. Monterey, CA: Brooks/Cole Publishing Company, 1975.
- [10] L. Austin, “Privacy and the Question of Technology,” *Law and Philosophy*, vol. 22, no. 2, pp. 119–166, 2003.
- [11] B.-J. Koops, B. C. Newell, T. Timan, I. Skorvanek, T. Chokrevski, and M. Galic, “A typology of privacy,” *U. Pa. J. Int’l L.*, vol. 38, p. 483, 2016.
- [12] D. S. Syrdal, M. L. Walters, N. Otero, K. L. Koay, and K. Dautenhahn, “He knows when you are sleeping—privacy and the personal robot companion,” in *Workshop on the Human Implications of Human-Robot Interaction, Association for the Advancement of Artificial Intelligence (AAAI ’07)*, 2007, pp. 28–33.
- [13] K. Caine, S. Sabanovic, and M. Carter, “The Effect of Monitoring by Cameras and Robots on the Privacy Enhancing Behaviors of Older Adults,” in *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI ’12. New York, NY, USA: ACM, 2012, pp. 343–350.
- [14] R. Y. Wong and D. K. Mulligan, “These aren’t the autonomous drones you’re looking for: investigating privacy concerns through concept videos,” *Journal of Human-Robot Interaction*, vol. 5, no. 3, pp. 26–54, 2016.
- [15] T. Denning, C. Matuszek, K. Koscher, J. R. Smith, and T. Kohno, “A spotlight on security and privacy risks with future household robots: attacks and lessons,” in *Proceedings of the 11th international conference on Ubiquitous computing*. ACM, 2009, pp. 105–114.
- [16] M. K. Lee, K. P. Tang, J. Forlizzi, and S. Kiesler, “Understanding users’ perception of privacy in human-robot interaction,” in *6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2011, pp. 181–182.
- [17] A. Hubers, E. Andrulis, L. Scott, T. Stirrat, R. Zhang, R. Sowell, M. Rueben, C. M. Grimm, and W. D. Smart, “Using video manipulation to protect privacy in remote presence systems,” in *International Conference on Social Robotics*. Springer, 2015, pp. 245–254.

- [18] D. Butler, J. Huang, F. Roesner, and M. Cakmak, "The Privacy-Utility Tradeoff for Remotely Teleoperated Robots," in *Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, Portland, OR, 2015.
- [19] J. Klow, J. Proby, M. Rueben, R. T. Sowell, C. M. Grimm, and W. D. Smart, "Privacy, utility, and cognitive load in remote presence systems," in *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 2017, pp. 167–168.
- [20] M. Rueben, F. J. Bernieri, C. M. Grimm, and W. D. Smart, "Evaluation of physical marker interfaces for protecting visual privacy from mobile robots," in *Proceedings of the 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 2016)*. IEEE, 2016, pp. 787–794.
- [21] A. R. Wagner, "An autonomous architecture that protects the right to privacy," in *Proceedings of AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*. AAAI, 2018.
- [22] M. Rueben, F. J. Bernieri, C. M. Grimm, and W. D. Smart, "Framing effects on privacy concerns about a home telepresence robot," in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 2017, pp. 435–444.
- [23] M. Tonkin, J. Vitale, S. Ojha, J. Clark, S. Pfeiffer, W. Judge, X. Wang, and M.-A. Williams, "Embodiment, privacy and social robots: May I remember you?" in *International Conference on Social Robotics*. Springer, 2017, pp. 506–515.
- [24] J. Vitale, M. Tonkin, S. Herse, S. Ojha, J. Clark, M.-A. Williams, X. Wang, and W. Judge, "Be more transparent and users will like you: A robot privacy and user experience design experiment," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 2018, pp. 379–387.
- [25] M. M. Krupp, M. Rueben, C. M. Grimm, and W. D. Smart, "A focus group study of privacy concerns about telepresence robots," in *Proceedings of the 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 2017)*. IEEE, 2017, in review.
- [26] M. R. Calo, "The drone as a privacy catalyst," *Stan. L. Rev. Online*, vol. 64, p. 29, 2011.
- [27] M. E. Kaminski, "Robots in the Home: What Will We Have Agreed To?" 2015.
- [28] M. E. Kaminski, M. Rueben, W. D. Smart, and C. M. Grimm, "Averting robot eyes," *Md. L. Rev.*, vol. 76, p. 983, 2016.
- [29] C. Lutz and A. Tamò, "Robocode-ethicists: Privacy-friendly robots, an ethical responsibility of engineers?" in *Proceedings of the ACM Web Science Conference*. ACM, 2015, p. 21.
- [30] —, "Privacy and healthcare robots—an ANT analysis," in *We Robot 2016: the Fifth Annual Conference on Legal and Policy Issues relating to Robotics*. University of Miami School of Law, 2016, Discussant: Matt Beane, University of California Santa Barbara.
- [31] T. Schulz and J. Herstad, "Walking away from the robot: Negotiating privacy with a robot," in *Proceedings of British HCI Conference (BISL)*, 2017, pp. 1–6.
- [32] E. Sedenberg, J. Chuang, and D. Mulligan, "Designing commercial therapeutic robots for privacy preserving systems and ethical research practices within the home," *International Journal of Social Robotics*, vol. 8, no. 4, pp. 575–587, 2016.
- [33] M. Rueben and W. D. Smart, "Privacy in human-robot interaction: Survey and future work," in *We Robot 2016: the Fifth Annual Conference on Legal and Policy Issues relating to Robotics*. University of Miami School of Law, 2016, Discussant: Ashkan Soltani, Independent Researcher.
- [34] G. Hu, W. P. Tay, and Y. Wen, "Cloud robotics: architecture, challenges and applications," *IEEE network*, vol. 26, no. 3, 2012.
- [35] J. John J. Trinckes, *The Definitive Guide to Complying with the HIPAA/HITECH Privacy and Security Rules*. Auerbach Publications, 2012.
- [36] S. Booth, J. Tompkin, H. Pfister, J. Waldo, K. Gajos, and R. Nagpal, "Piggybacking Robots: Human-Robot Overtrust in University Dormitory Security," in *Proceedings of the twelfth annual ACM/IEEE international conference on Human-Robot Interaction - HRI '17*, 2017, pp. 426–434.
- [37] S. Gibbs, "Hackers can hijack Wi-Fi Hello Barbie to spy on your children," 2015. [Online]. Available: <https://www.theguardian.com/technology/2015/nov/26/hackers-can-hijack-wi-fi-hello-barbie-to-spy-on-your-children>
- [38] L. Franceschi-Bicchierai, "How This Internet of Things Stuffed Animal Can Be Remotely Turned Into a Spy Device," 2017. [Online]. Available: https://motherboard.vice.com/en_us/article/qkm48b/how-this-internet-of-things-teddy-bear-can-be-remotely-turned-into-a-spy-device
- [39] A. Cavoukian, *Privacy by Design: The 7 Foundational Principles*. Information and Privacy Commissioner of Ontario, 2016, <http://www.privacybydesign.ca/index.php/about-pbd/7-foundational-principles/>.
- [40] W. N. Hartzog, "Et Tu, Android? Regulating Dangerous and Dishonest Robots," *Journal of Human-Robot Interaction*, vol. 5, no. 3, p. 70, 2016.
- [41] N. Richards and W. Hartzog, "Taking trust seriously in privacy law," *Stan. Tech. L. Rev.*, vol. 19, p. 431, 2015.
- [42] D. M. Rousseau, S. B. Sitkin, R. S. Burt, and C. Camerer, "Not so different after all: A cross-discipline view of trust," *Academy of management review*, vol. 23, no. 3, pp. 393–404, 1998.
- [43] D. H. McKnight, V. Choudhury, and C. Kacmar, "Developing and validating trust measures for e-commerce: An integrative typology," *Information systems research*, vol. 13, no. 3, pp. 334–359, 2002.
- [44] P. A. Hancock, D. R. Billings, K. E. Schaefer, J. Y. C. Chen, E. J. de Visser, and R. Parasuraman, "A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 53, no. 5, pp. 517–527, 2011.
- [45] E. Buyuksagis and H. V. B. Willem, "Strict liability in European codification. Torn between objects, activities and their risks," *Geo. J. Int'l L.* 44, p. 609, 2012.
- [46] H. Farber, "To testify or not to testify: A comparative analysis of Australian and American approaches to a parent-child testimonial exemption," *Tex. Int'l LJ*, vol. 46, p. 109, 2010.
- [47] D. J. Solove and W. Hartzog, "The FTC and the new common law of privacy," *Colum. L. Rev.*, vol. 114, p. 583, 2014.
- [48] D. Simshaw, N. Terry, K. Hauser, and M. Cummings, "Regulating healthcare robots: Maximizing opportunities while minimizing risks," *Rich. JL & Tech.*, vol. 22, p. 1, 2015.
- [49] S. Barth and M. D. T. De Jong, "The Privacy Paradox—Investigating Discrepancies between Expressed Privacy Concerns and Actual Online Behavior—A Systematic Literature Review," *Telematics and Informatics*, vol. 34, pp. 1038–1058, 2017.
- [50] L. Baruh, E. Secinti, and Z. Cemalcilar, "Online Privacy Concerns and Privacy Management: A Meta-Analytical Review," *Journal of Communication*, vol. 67, no. 1, pp. 26–53, 2017.
- [51] S. Kokolakis, "Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon," *Computers & Security*, vol. 64, pp. 122–134, 2017.
- [52] H. J. Smith, T. Dinev, and H. Xu, "Information privacy research: an interdisciplinary review," *MIS quarterly*, vol. 35, no. 4, pp. 989–1016, 2011.
- [53] A. E. Marwick and D. Boyd, "Networked privacy: How teenagers negotiate context in social media," *New Media & Society*, vol. 16, no. 7, pp. 1051–1067, 2014.
- [54] T. Dinev and P. Hart, "An extended privacy calculus model for e-commerce transactions," *Information Systems Research*, vol. 17, no. 1, pp. 61–80, 2006.
- [55] J. E. Campbell and M. Carlson, "Panopticon.com: Online Surveillance and the Commodification of Privacy," *Journal of Broadcasting & Electronic Media*, vol. 46, no. 4, pp. 586–606, 2002.
- [56] A. R. Beresford, D. Kübler, and S. Preibusch, "Unwillingness to pay for privacy: A field experiment," *Economics Letters*, vol. 117, no. 1, pp. 25–27, 2012.
- [57] D. J. Solove, *Understanding privacy*. Cambridge, MA: Harvard University Press, 2008.