D. Reinhardt and A. Misra. Toward Privacy-Respecting Service Robots. IEEE Pervasive Computing, 2025. DOI: 10.1109/MPRV.2025.3555407

Article Type: Feature Article

Towards Privacy-respecting Service Robots

Delphine Reinhardt, *University of Göttingen, Goldschmidtstr. 7, Göttingen, Germany*

Archan Misra, *Singapore Management University, 80 Stamford Road, Singapore, Singapore*

Abstract—Future service robots have the potential to improve our quality of life across diverse contexts. However, they will introduce multi-faceted privacy threats. After detailing such threats, we outline possible research approaches to elicit and implement user privacy preferences, considering robots' capabilities and their deployment contexts, including multi-user scenarios in public spaces. Our work aims to catalyze future research efforts to promote the responsible, privacy-compliant deployment of such robots.

A new era of service robots is emerging, driven by advancements in GenAI, humanoid robotics, natural language and multi-modal reasoning. Such robots will be characterized by two capabilities: (a) the ability to derive semantic understanding of our surrounding physical environments, creating both greater autonomy in navigation and enhanced situational awareness, and (b) the ability to engage more naturally with humans via bidirectional multi-modal interfaces, including visual, speech and gestural cues. Robots embedded with such AI-based interaction and reasoning capabilities will be deployed beyond traditional industrial assembly lines, in diverse environments such as homes, museums, and hospitals.

To realize such advanced reasoning and interaction capabilities, the robots will increasingly be equipped with a variety of sensors that collectively provide fine-grained 3D sensing of their environment and the ability to comprehend human commands and queries. Their proliferation will introduce new forms of, or elevate existing, privacy threats that go beyond digital information captured by current devices.

To promote responsible deployment and mitigate consumer concerns, we identify the multi-faceted privacy threats and propose a research agenda to address such threats through new capabilities for eliciting and applying user's privacy preferences. A major issue is that existing solutions to capture individual privacy preferences, such as manually selecting from a predefined list of alternatives, will become even more impractical as more privacy dimensions in more contexts with more bystanders will be prevalent. We hence argue for two main approaches: (1) Robots will learn explicit and implicit privacy preferences based on natural behavioral markers. To foster transparent communication of such multi-faceted privacy risks, these preferences will be made available to the users on their personal device, so that they can also be transferred, applied, and updated during encounters with other robots, (2) The privacy-preserving behavior of robots will combine such learned user preferences with a deeper understanding of the environmental and social context, including the presence of multiple people [1], the sensitivity of the current task, and privacy vs. utility vs. safety tradeoffs.

Among others, these approaches are aligned with the privacy-by-design theory [2], the concept of user-centrism also supported in [3] and other privacy principles especially considered in [4] including the principle of data minimization further required by the European *General Data Protection Regulation* (GDPR) (Art. 5.1.(c)). By discussing these approaches and related pitfalls, we outline directions for the research community to pursue, to better protect privacy and foster trust in a "robot-saturated" future.

MULTI-FACETED PRIVACY THREATS

While improving our quality of life, service robots raise privacy threats common to related applications, but also generate new ones as illustrated in Tab. 1. In a nutshell, they can be deployed for a variety of tasks, spanning different locations and exhibiting multiple forms of robot-human relationships (one-to-one, one-to-many and many-to-many). The robots need to (1) analyze and understand their environment, (2) observe multi-modal interactions between one or more humans,

XXXX-XXX © 2025 IEEE
Digital Object Identifier 10.1109/MPRV.2025.3555407

TABLE 1: Comparison of privacy threats and dimensions (I: Informational, S: Social, P: Physical, B: Bodily) between robots and related applications.

	Cameras	Microphones	Navigation	Mobility	Communication				Inter-agent collaboration	Deployment spaces		
					Speech	Display	Gestures	Proactivity		Private	Semi-public	Public
Privacy dimensions	I	I,S	I	I,S,P,B	S	I	P,B	I,S,P,B	I			
Surveillance cameras [5]	x									x	x	x
Smart speakers [6], [7]		x			x					x		
Smart home devices [8], [9]										x		
Service robots	x	x	x	x	x	x	x	x	x	x	x	x

and (3) interact with humans with increasing proactivity and through different modalities (verbal, gestural, visual and tactile). Moreover, they may not only infringe users' informational privacy like most existing applications, but also users' social, physical, and bodily privacy [10], [11]. For example, users may have different preferences in terms of data being collected [12], [10] (informational privacy). Some users may prefer robots that only react to their commands, while others may enjoy more "chatty" robots that engage proactively, depending on their own personality [13] (social privacy). Other users may prefer that robots showing human faces may operate at a closer distance [14] (physical privacy), while users may have different preferences about the way robots touch them [15] (bodily privacy). We now detail the privacy threats that arise from a set of common robotic sensing and interaction modalities following the same categories as in Tab. 1.

Cameras

Classified among the most privacy-invasive sensors [16], cameras on robots can capture sensitive information and personal activities, such as personal hygiene [17]. Capturing users' gestures and faces can reveal information about their emotions and health status. Face detection can further capture the number of people present and identify them using online databases. Objects observed in private spaces can reveal information about personality (untidiness) and tastes (colorful/monochromatic decoration). As many personal spaces are also shared (e.g., the living room in an apartment), the camera can effectively capture such attributes not just for the primary user, but also for other occupants.

Microphones

Robots are embedded with microphones that are continuously active to capture and process users' voice commands. They not only capture the speech of their owners, but those of bystanders, too. As another highly privacy-invasive sensor [18], microphones can not only identify the speaker, but the captured voice data can also reveal other attributes, e.g., health status [19]. Background noises, such as traffic, provide further information about the environment. Although these risks are similar to those of current smart speakers, capturing image data concurrently is likely to exacerbate privacy concerns.

Navigation

To support autonomous navigation, robots are often equipped with radars and LIDARs, which provide continuous 3D sensing of their environment. While such 'non-visual' sensors are often viewed as more privacy-friendly than cameras, sensed 3D point cloud data can also be used to extract sensitive information, such as

facial expressions [20].

Mobility

Being mobile, the robot might go into places where users would prefer to be left alone, such as the bathroom [10], breaching users' social privacy. Approaching too closely can violate users' physical privacy by invading their "personal space" and touching them may violate their bodily privacy.

Communication

Robots will leverage AI advances in vision, natural language processing, and wireless gesture recognition for more natural and immersive voice- and gesture-based interactions. They are likely to respond to user instructions and queries via natural language responses, through synthesized speech, on the robot's display, and/or via gestures. Because these interaction modalities are less private, it is difficult for users to control what information is made available to whom in their vicinity. This concern is further heightened by the inevitable personalization and proactivity that such agents will possess. Similar to smart speakers, the robots may reveal more information to other family members or bystanders than the subjects would like. For example, such agents can reveal information, either in reaction to user queries or proactively based on inferred context, that may be inappropriate for the current users' social context. The key issue is that the robot may not understand the social contexts, potentially leading to an inappropriate behavior [1].

Inter-Agent Collaboration

Until now, we have only considered a single robot interacting with one or multiple individuals. However, in the near future, we may witness a more collaborative scenario, where multiple robots, each tuned to a specialized task, may mutually exchange information and complement their respective actions to reach a common goal. As a result, they may have more capabilities when working together and gain access to additional richer information about the users and leverage it in cross-contexts, leading to problematic outcomes for some users [1].

Deployment Spaces

As compared to private deployment, robots deployed in semi-public to public spaces exacerbate the threats to privacy [21]. Instead of being acquired, configured, and deployed at home by the users themselves like most existing technologies (see Tab. 1), future robots will span a wider range of locations and greater population segments with less control on privacy protection. Instead of only disclosing information to visitors at home, robots may also reveal sensitive information to a wider range of bystanders in public spaces. Threats to physical and bodily privacy are especially relevant in these spaces, as the robots are not personal and users may feel that they are being accosted by a stranger.

ADDRESSING THREATS

To address these threats, guidelines and recommendations have been proposed based on conceptual models, such as Palen and Dourish's framework in [22], the Fair Information Practice Principles in [23], and the privacy-by-design principle [24] based on [2]. In addition to transparency requirements [25], further recommendations for explicit [11], dynamic [23], [26], or cascade consent [27] have been made. Different technical solutions exist that partly support these suggestions. They range from trivial (e.g., privacy notices like light signals [28] (perceived as insufficient in [18]) or relying on more privacy-friendly sensors [29]) to more complex techniques (e.g., hiding sensitive information [30]). In contrast, we consider how users' preferences can be both elicited and applied along the life cycle of their encounters with robots. Fig. 1 compiles the related needs for research that we discuss in what follows.

ELICITING PRIVACY PREFERENCES

Given the diversity of environments and contexts within which robots will be deployed, a reductive one-size-fits-all, regulatory approach seems infeasible. Instead, we require a context-aware approach that can reflect users' and bystanders' privacy preferences and regulate the corresponding human-robot interactions to respect their boundaries along the privacy dimensions highlighted in Tab. 1. Clearly, these preferences are also likely to vary with users' context and moods. For example, users may be more accepting of privacy-invasive data collection in cases of emergency, of interruptions when unoccupied, and of robots approaching closer when handing over objects. We believe that new research is needed to develop both the principles and mechanisms by which such a diverse and context-dependent set of privacy preferences can be addressed in a practical manner. Our articulated research directions can help to (1) gather such context-dependent, individualized preferences in a user-friendly fashion as recommended by, e.g., [2], [3], and (2) reflect them in appropriate behavioral adaptation by the robots.

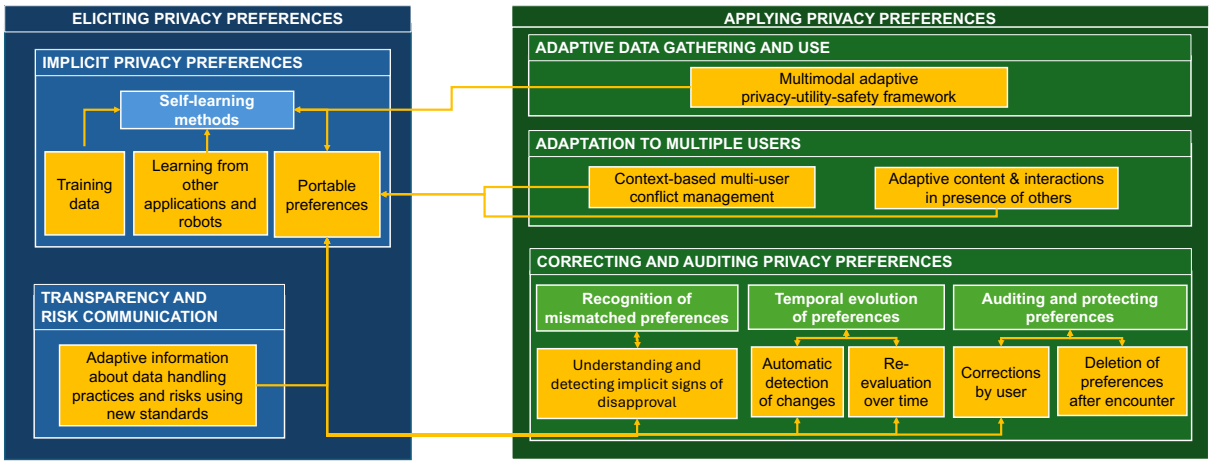

FIGURE 1: Overview of the different categories of proposed approaches and their relationships including the currently missing components (in yellow)

Implicit Privacy Preferences

For decades, users have been asked to manually set their privacy preferences in different contexts, such as content sharing in online social media or granting access to mobile applications. This approach is, however, inadequate for social robots because (1) users first need to translate and express their own privacy preferences into privacy settings [31] and (2) appropriate techniques are needed to implement such settings.

Robots are expected to be highly diverse in form, sensing and actuation capabilities. Also, it is highly unlikely that users will be able to explicitly establish their privacy preferences when interacting with such robots in public (and even semi-public) locations. Moreover, users should not only express their preferences about data collection and sharing as experienced today in, e.g., online social media, but also manage the boundaries of the other privacy dimensions (i.e., social, physical, bodily). It will hence be impractical to utilize existing approaches for manual, explicit setting of privacy preferences via graphical interfaces.

Instead, we need to develop novel solutions that assist users in expressing context-dependent privacy preferences. To this end, we propose to leverage (1) self-learning methods to reduce the configuration overhead for the users to the minimum, and (2) robotic sensing and interaction modalities. Besides explicit context (e.g., location, task) considered in existing solutions, implicit context can also be incorporated. This means using human behavior signals to more accurately map natural preferences according to societal standards. For example, a user desiring some physical privacy may simply lower their voice, wave their hands or leave the room as observed in [17], to indicate that the robot should leave them alone. These preferences can also vary across individuals and cultures. While exploiting such implicit behavioral cues, we must avoid the pitfall of failing to consider cultural differences or environmental context. There should also be a low-overhead mechanism to transfer such “privacy profiles" to new robots. This is especially important in semi-private and public spaces, where users do not have administrative ownership of the devices and human-machine interactions are often transient.

Self-learning Methods ML techniques are increasingly applied to predict sensitivity [31], [30] and user preferences [32]. While existing results, are promising, we lack not only relevant models but also the relevant training data for physical environments where human-robot collaboration will be commonplace. Training data are crucial; without them, learning may take too long and frustrate users. They must also cover various contexts, environments, and cultures. One potential approach is using recent advancements in natural scene and language understanding (via LLM and VLM models) to extract natural semantic descriptions of scenes (e.g., “a mother is talking to her child") and train AI models to predict privacy preferences given such semantic inputs, thus going beyond the identification of sensitive objects only (e.g., [30]). With such trained models, we may be able to bootstrap deployments successfully, with a majority of the predicted preferences matching users’ actual preferences. However, because of variability across users and the likely evolution in individual preferences, we will still need to embed incremental learning techniques in such models.

Learning from Other Applications and Robots Another approach to bootstrap the learning process is to carefully utilize results obtained in orthogonal domains. While they do not cover all relevant privacy dimensions, they provide useful exemplars for issues related to informational privacy. For example, existing text-based and image-based taxonomies initially developed (e.g., in [33]) and applied (e.g., in [31]) for online scenarios can be leveraged to identify sensitive conversations or observed scenes. Similarly, we should investigate if, and to what extent, privacy preferences for other widely-deployed technologies, such as mobile phones or smart speakers, may be transferred to robots. For example, a user who has not activated her voice assistant on her phone may be less willing to have her voice recorded by robots. Over-generalizing such preferences, without considering the current context and environment, is however a pitfall to be avoided.

Portable Privacy Preferences Another approach worth investigating is the use of personal mobile devices to serve as a privacy preference repository to bootstrap the robots' behavior. Instead of requiring an individual to undertake a bootstrapping phase with each newly encountered robot or having a user preferences being stored in a centralized repository, we can imagine that individual preferences are sent to a nearby robot using short-range wireless techniques as proposed in another context [34]. This approach can be especially useful in semi-private and public spaces where the individual-robot interaction may be transient or intermittent. However, an important challenge is the need to further adapt the robot's privacy-related behavior to the current interaction context. For example, a hospital concierge robot may avoid using its voice interface if the patient's condition is sensitive. We believe that the development of appropriate interactions to support such flexible, low-overhead privacy adaptation should be an area of active research.

Transparency and Risk Communication

In many cases, the user's privacy choice is not absolute but likely to be influenced by variations in how the collected data is used (e.g., who can access it and for what purpose) by the robots. To enable informed and individualized decision making, the robot needs to inform users on what data are collected and how such data are processed, stored, and disclosed. Moreover, being transparent about the privacy settings can influence both users' self-disclosure willingness and depth [35]. Such transparency capabilities are not only recommended [27] and desired by potential users [10], but also required by the GDPR in a concise, accessible, and understandable form (Recital 58). However, to our knowledge, no solutions taking into account the diversity of robots, their interaction capabilities and their deployment scenarios exist.

Designing such solutions is, however, fraught with potential pitfalls. Indeed, determining information to be provided is a nontrivial task [25], because it depends on the type of robot and the varying information needs of individuals [36]. Creating methods to explain the tradeoff between benefits, risks (including safety concerns) while using certain sensors is crucial. This helps users find a tradeoff between worrying excessively and being unaware of privacy risks. Moreover, both the content and form of communication matters. For example, displays may not be available on robots. One promising approach is leveraging LLMs to generate and orally communicate easy-to-comprehend summaries, and possibly provide interactive clarifications to specific situations. For a consistent user experience, we also need appropriate standards for related information content, format, and interaction patterns.

Research recommendation: We need to develop new solutions to infer and learn user preferences implicitly (as opposed to relying only on explicit manual input) by leveraging self-learning models, cross-domain knowledge, and context-aware adaptations. To ensure privacy compliance even during transient human-robot interactions, we should also explore the seamless transfer of privacy preferences stored on personal devices. Finally, we need to develop techniques to communicate privacy risks transparently and interactively, aided by standards and the use of LLMs.

APPLYING PRIVACY PREFERENCES

We now consider the equally complex topic of how robots should adapt their behavior to user preferences.

Adaptive Data Gathering and Use

Even if we assume that the robots can be trusted to respect individual privacy preferences, software developers still face challenges in programmatically ensuring that robotic behavior complies with the desired outcomes. Even in the limited context on informational privacy for existing mobile devices, developers already face design and implementation challenges [37]. To avoid existing pitfalls, developers that are not expert in privacy will therefore need more support, as respecting user preferences will require them to consider a larger set of attributes and values.

Such support will include developing a better un-

derstanding of the tradeoff between losses in utility, increase in privacy, and safety guarantees caused by restrictive data sensing and collection preferences. For example, it is likely that a robot can reduce the spatial resolution and frequency of LIDAR-based sensing; in turn, this may imply lower precision in its localization and navigation capabilities. Such reduced resolution within an adaptive LIDAR scanning mechanism may be sufficient for sensing the aggregate structure of, e.g., a kitchen, with a fine-grained resolution being needed only to provide cooking instructions. A pitfall of reduced navigation precision is increased safety risks in sensitive environments, such as a robot transporting hot liquids near a child. While proxemics studies (e.g., [38]) explore social interaction preferences with robots, more research is needed to determine how spatial accuracy should vary based on tasks and environments, and when safety concerns justify overriding human privacy preferences.

To develop such an adaptive framework, we first need to build a catalog of task contexts and the relationship between the required accuracy/fidelity of such tasks and the parameters of the underlying sensors. Such a catalog, by necessity, must be multi-modal: for example, the accuracy of object recognition may need to be characterized in terms of the resolution of both RGB vision and 3D LIDAR scanners and also depend on the environmental context, with 3D LIDAR proving more informative than RGB under low-lighting conditions. Such framework will conform to the concept of data minimization requested by the GDPR (Art. 5.1.(c)).

Adaptation to Multiple Users

The presence of additional users in the environment generate additional challenges for privacy-compliant human-robot interaction.

Context-based Multi-User Conflict Management Assume that a father and his young child are interacting with a hospital concierge robot. Both may possess conflicting preferences. For example, the visually impaired father might prefer the robot to come closer for easier display reading, while the child might prefer it to stay farther away due to intimidation. “Obvious" approaches to handle such conflicts, such as (a) having the robot continually adjust its distance based on whether it interacts with the father or child or (b) always defaulting to the most conservative preference, are not satisfactory. Solutions should hence be developed to negotiate these conflicts, based, e.g., on identified social relationships as considered in [1], and to find context-aware solutions that optimize collective utility metrics, e.g., prioritizing the responsible adult’s preference in this case. A pitfall, however, is that such conflict negotiation techniques may need to be re-executed for each new user entering the scene.

Adaptive Content & Interactions in Presence of Others As noted earlier, robots that utilize publicly observable communication channels may disclose unintended content, potentially causing embarrassment. We hence need privacy-conscious solutions that will adapt the sensitivity of the disclosed information and the nature of interactions to the relevant environmental and social context. For example, a robot can remind a user using a subtle verbal hint or use a discreet channel (e.g., a push notification on the user’s device) in the presence of visitors. It will, however, be challenging to identify and define the strength of social relationships without preliminary knowledge, especially in public environments. Such adaptation may also extend to the robot’s own physical actions. For example, a robot being asked to “prepare us a drink" may observe that a pregnant friend is present, and may thus not consider an alcoholic drink. We believe that this general area, of interaction content, interaction modality, physical action adaptation to broader environmental context, is a novel dimension of privacy, essential for human-robot interaction, that requires further research.

Correcting and Auditing Privacy Preferences

Recognition of Mismatched Preferences The robots’ reactions and interactions may not be fully aligned with the user preferences, especially in the initial bootstrapping phase. Developing mechanisms that both continually detect user dissatisfaction with such robotic behaviors and then update user privacy profiles and robot behavioral outcomes is thus extremely critical. Ideally, the robot should first autonomously perceive that its behavior was not appropriate by capturing the corresponding users’ cues, either explicitly (e.g., speech preferred by most users in [39]) or implicitly (e.g., raised eyebrows). While perception of such cues is an intrinsic part of human-human social interaction, we currently do not know how users would react in human-robot interaction contexts–i.e., which cues would humans commonly utilize when interacting with inanimate machines to convey privacy concerns? We can speculate that explicit cues may replicate the common privacy-preserving behaviors adopted when interacting with humans, such as asking the robot to stay silent. However, the transferrability of implicit cues is currently largely unknown, both in terms of what humans would prefer to use and also what implicit cues can be reliably sensed by the robots. Also, it

is likely that the preferred implicit cues will also vary with differences in robot attributes (e.g., size, speed of motion), user demographics and environmental/task contexts. Accordingly, we suggest an active research agenda around both (a) understanding implicit signs of users' disapproval in different contexts, and (b) evaluating the technical feasibility of detecting them using the robot's sensors.

Temporal Evolution of Preferences Another pitfall is that user preferences are likely to change over time, both temporarily and permanently. For example, users living alone with a robot may have a set of privacy preferences that dramatically change after key life events (e.g., new cohabitation with partner). Similarly, users returning home after work may be exhausted on certain days, and prefer that a robot is less chatty and leaves them alone that evening. In these situations, an individual's intrinsic state may have direct, but transient, consequences on both the social and physical boundaries that robots should respect.

We further anticipate that the user preferences will evolve with their experience with such robots. We can expect users to first be more conservative, but then gradually become more permissive as they get familiar with robots deployed in different contexts. As often occurs in human-to-human relationships, users may progressively socially bond with the robots, thus potentially willing to reveal more information, allowing them to interrupt them more when speaking, or approaching them closer. The extent to which such user-robot relationships may suffer, in terms of both gravity and duration, depending on the context is unclear; we note that utility and entertainment benefits paradoxically outweighed a simulated privacy breach in an online clothes shopping study [40]. Instead of pop-ups to ask users to verify their preferences, we hence suggest creating solutions that automatically detect changes by picking up relevant user cues. Moreover, user preferences should be re-evaluated as the technology evolves in terms of, e.g., new features or possible privacy breaches.

Auditing and Protecting Preferences To foster trust in the robots, we encourage making the learned preferences transparent to the users. This will allow users to inspect such preferences, audit them, and manually correct them if necessary. These privacy preferences themselves may be considered as sensitive personal information, thereby subject to regulations and requiring appropriate care in exposing them only to authorized users. To prevent such privacy breaches, we need to develop mechanisms that allow a genuine update of user parameters by the robot based on its observation during the interactions, but prevent their exploitation and storage beyond the actual interaction episodes. In other words, the robots should forget about user preferences when their interactions with the robot end, especially in public environments.

Research recommendation: Robots should dynamically adapt data collection and interactions based on the tasks, the social and environmental context, and user preferences while abstaining from using intrusive methods like pop-ups. As human-robot relationships evolve, the robots must detect and adjust to changes in preferences, accommodating both increased familiarity and potential setbacks. Users should be able to audit and correct their preferences, which should be especially protected and forgotten after the interactions.

CONCLUSION

We have highlighted novel privacy challenges from the growing deployment of service robots. Their proactive engagement with humans through verbal, gestural, navigational, and tactile modalities demands solutions that address not only informational but also social, physical, and bodily privacy across diverse contexts. Among others, we emphasized the need for robots to implicitly learn privacy preferences and transparently communicate the risks and benefits of collecting sensitive data beyond traditional interfaces. This calls for a multi-disciplinary research agenda to develop accessible, privacy-preserving solutions for human-robot collaboration.

ACKNOWLEDGMENTS

We thank the anonymous reviewers, and Profs. Anthony Tang and Jiannan Li at SMU for relevant discussions. Archan Misra's work was supported by National Research Foundation, Prime Minister's Office, Singapore, both under its NRF Investigatorship grant (NRF-NRFI05-2019-0007), and The Mens, Manus, and Machina (M3S) interdisciplinary research group (IRG) under its Campus for Research Excellence and Technological Enterprise (CREATE) program.

REFERENCES

1. S. Reig, M. Luria, J. Z. Wang, D. Oltman, E. J. Carter, A. Steinfeld, J. Forlizzi, and J. Zimmerman, "Not Some Random Agent: Multi-person Interaction with a Personalizing Service Robot," in *Proc. ACM/IEEE HRI*, 2020.

2. A. Cavoukian *et al.*, “Privacy by Design: The 7 Foundational Principles,” *Information and Privacy Commissioner of Ontario, Canada*, 2009.
3. I. S. Rubinstein and N. Good, “Privacy by Design: A Counterfactual Analysis of Google and Facebook Privacy Incidents,” *Berkeley Tech. LJ*, 2013.
4. A. Chatzimichali, R. Harrison, and D. Chrysostomou, “Toward Privacy-sensitive Human–robot Interaction: Privacy Terms and Human–data Interaction in the Personal Robot Era,” *Paladyn, Journal of Behavioral Robotics*, 2020.
5. S. Lederer, J. I. Hong, A. K. Dey, and J. A. Landay, “Personal Privacy through Understanding and Action: Five Pitfalls for Designers,” *Personal and Ubiquitous Computing*, 2004.
6. J. Lau, B. Zimmerman, and F. Schaub, “Alexa, are you listening? Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers,” *Proc. ACM on HCI*, 2018.
7. N. Malkin, J. Deatrick, A. Tong, P. Wijesekera, S. Egelman, and D. Wagner, “Privacy Attitudes of Smart Speaker Users,” *PoPETs*, 2019.
8. S. Zheng, N. Apthorpe, M. Chetty, and N. Feamster, “User Perceptions of Smart Home IoT Privacy,” *Proc. ACM on HCI*, 2018.
9. C. Geeng and F. Roesner, “Who’s in control? Interactions in Multi-user Smart Homes,” in *Proc. SIGCHI CHI*, 2019.
10. D. Reinhardt, M. Khurana, and L. H. Acosta, ““I still need my privacy”: Exploring the Level of Comfort and Privacy Preferences of German-speaking Older Adults in the Case of Mobile Assistant Robots,” *PMC*, 2021.
11. M. Rueben, A. M. Aroyo, C. Lutz, J. Schmölz, P. Van Cleynenbreugel, A. Corti, S. Agrawal, and W. D. Smart, “Themes and Research Directions in Privacy-Sensitive Robotics,” in *Proc. Workshop on Advanced Robotics and Its Social Impacts*, 2018.
12. C. Berridge and T. F. Wetle, “Why Older Adults and their Children Disagree about In-home Surveillance Technology, Sensors, and Tracking,” *The Gerontologist*, 2020.
13. S. Whittaker, Y. Rogers, E. Petrovskaya, and H. Zhuang, “Designing Personas for Expressive Robots: Personality in the New Breed of Moving, Speaking, and Colorful Social Home Robots,” *Journal HRI*, 2021.
14. N. E. Neef, S. Zabel, M. Lauckner, and S. Otto, “What is Appropriate? On the Assessment of Human-robot Proxemics for Casual Encounters in Closed Environments,” *International Journal of Social Robotics*, 2023.
15. A. Hitzmann, H. Sumioka, and M. Shiomi, “Touch Me Right: Lateral Preferences During Touch in Human-Robot-Interactions,” in *Proc. IEEE RO-MAN*, 2023.
16. F. J. R. Lera, C. F. Llamas, Á. M. Guerrero, and V. M. Olivera, “Cybersecurity of Robotics and Autonomous Systems: Privacy and Safety,” *Robotics-Legal, Ethical and Socioeconomic Impacts*, 2017.
17. K. Caine, S. Šabanović, and M. Carter, “The Effect of Monitoring by Cameras and Robots on the Privacy Enhancing Behaviors of Older Adults,” in *Proc. ACM/IEEE HRI*, 2012.
18. I. Ahmad, R. Farzan, A. Kapadia, and A. J. Lee, “Tangible Privacy: Towards User-centric Sensor Designs for Bystander Privacy,” *Proc. ACM on HCI*, 2020.
19. L. H. Acosta and D. Reinhardt, “A Survey on Privacy Issues and Solutions for Voice-controlled Digital Assistants,” *Pervasive and Mobile Computing (PMC)*, 2022.
20. X. Zhang, Y. Zhang, Z. Shi, and T. Gu, “mm-FER: Millimetre-wave Radar based Facial Expression Recognition for Multimedia IoT Applications,” in *Proc. ACM MobiCom*, 2023.
21. 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 *Proc. ICSR*, 2017.
22. T. Schulz, J. Herstad, and H. Holone, “Privacy at Home: An Inquiry Into Sensors and Robots for the Stay at Home Elderly,” in *Proc. ITAP*, 2018.
23. 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*, 2016.
24. T. Heuer, I. Schiering, and R. Gerndt, “Privacy-Centered Design for Social Robots,” *Interaction Studies*, 2019.
25. H. Felzmann, E. Fosch-Villaronga, C. Lutz, and A. Tamo-Larrieux, “Robots and Transparency: The Multiple Dimensions of Transparency in the Context of Robot Technologies,” *Robotics & Automation Magazine*, 2019.
26. E. Fosch Villaronga, A. Tamò-Larrieux, and C. Lutz, “Did I Tell you my New Therapist is a Robot? Ethical, Legal, and Societal Issues of Healthcare and Therapeutic Robots,” *Ethical, Legal, & Societal Issues of Healthcare & Therapeutic Robots*, 2018.
27. N. Fronemann, K. Pollmann, and W. Loh, “Should my Robot Know what’s best for me? Human-Robot Interaction between User Experience and Ethical Design,” *AI & Society*, 2022.
28. P. Su and X. Yuan, “Are You Watching Me? A Study on Privacy Notice Design of Social Robot,” in *Proc. AHFE*, 2021.
29. S. Eick and A. I. Antón, “Enhancing Privacy in Robotics via Judicious Sensor Selection,” in *Proc. IEEE ICRA*, 2020.

30. F. E. Fernandes, G. Yang, H. M. Do, and W. Sheng, “Detection of Privacy-sensitive Situations for Social Robots in Smart Homes,” in *Proc. CASE*, 2016.
31. L. Kqiku and D. Reinhardt, “SensitivAlert: Image Sensitivity Prediction in Online Social Networks using Transformer-based Deep Learning Models,” in *Proc. AAAI ICWSM*, 2024.
32. T. Nakamura, S. Kiyomoto, W. B. Tesfay, and J. Serna, “Easing the Burden of Setting Privacy Preferences: A Machine Learning Approach,” in *Proc. ICISSP*, 2017.
33. A. J. Gill, A. Vasalou, C. Papoutsi, and A. N. Joinson, “Privacy Dictionary: A Linguistic Taxonomy of Privacy for Content Analysis,” in *Proc. CHI*, 2011.
34. N.-W. Gong, M. Laibowitz, and J. A. Paradiso, “Dynamic Privacy Management in Pervasive Sensor Networks,” in *Proc. AmI*, 2010.
35. J. Stapels, A. Augustine, N. Diekmann, and F. Eyssel, “Never Trust Anything That Can Think for Itself, if You Can’t Control Its Privacy Settings: The Influence of a Robot’s Privacy Settings on Users’ Attitudes and Willingness to Self-disclose,” *International Journal of Social Robotics*, 2023.
36. P. Murmann, D. Reinhardt, and S. Fischer-Hübner, “To Be, or Not To Be notified: Eliciting Privacy Notification Preferences for Online mHealth Services,” in *Proc. IFIP SEC*, 2019.
37. M. Tahaei, R. Abu-Salma, and A. Rashid, “Stuck in the Permissions with you: Developer & End-user Perspectives on App Permissions & their Privacy Ramifications,” in *Proc. CHI*, 2023.
38. P. Patompak, S. Jeong, I. Nilkhamhang, and N. Chong, “Learning Proxemics for Personalized Human–Robot Social Interaction,” *International Journal of Social Robotics*, 2020.
39. A. Austermann and S. Yamada, ““Good Robot”,“Bad Robot”—Analyzing Users’ Feedback in a Human-robot Teaching Task,” in *Proc. IEEE RO-MAN*, 2008.
40. G. Hannibal, A. Dobrosovestnova, and A. Weiss, “Tolerating Untrustworthy Robots: Studying Human Vulnerability Experience within a Privacy Scenario for Trust in Robots,” in *Proc. IEEE RO-MAN*, 2022.

Delphine Reinhardt is a full professor for Computer Security and Privacy at the University of Göttingen, Germany. She was a visiting professor at Singapore Management University (SMU) in 2023 and 2024. She completed her doctoral thesis in computer science at TU Darmstadt. Her research interests include usable privacy for emerging technologies.

Archan Misra is the Lee Kong Chian Professor of Computer Science and Vice Provost (Research) at SMU. He holds a Ph.D. in Electrical and Computer Engineering from the University of Maryland at College Park. He has worked extensively on problems spanning mobile & pervasive computing, wireless networking, and human-machine collaboration.