Rules for Robots: Co-Designing Privacy-Conscious Social Robots with Older Adults

Clara Y. Xi
clara.xi@ucalgary.ca
University of Calgary
Calgary, Alberta, Canada

Lora Oehlberg
lora.oehlberg@ucalgary.ca
University of Calgary
Calgary, Alberta, Canada

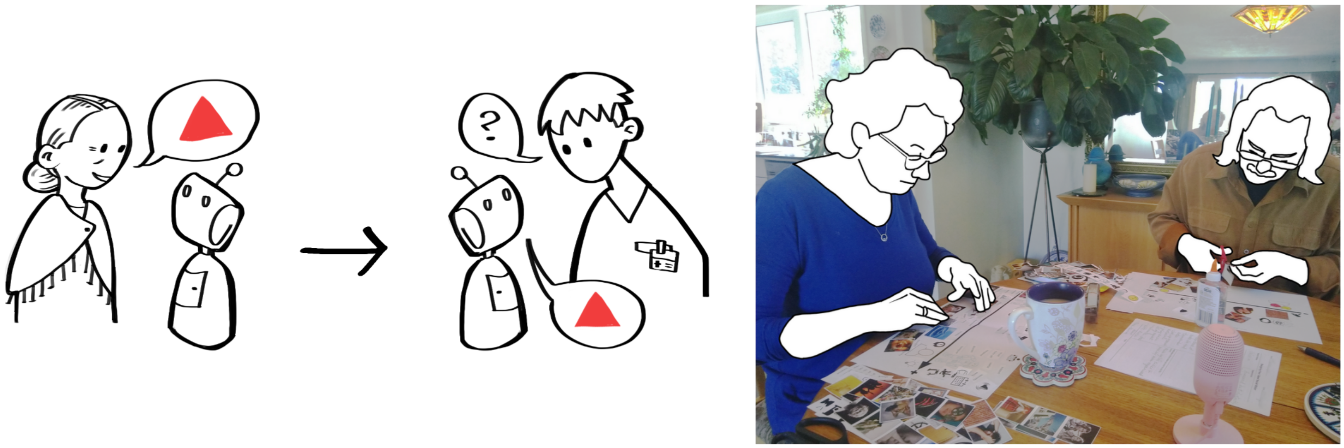

Figure 1: As social, embodied technologies, social robots are likely to be implicated in the interpersonal privacy of older adult users (see left). We present the results of a co-design study (see right) which investigated how social robots may be designed to protect this dimension of their users' privacy.

Abstract

While researchers project that social robots may one day be prevalent in the lives of older adults, privacy issues remain underexamined. Social robots can present unique and significant threats to their users' privacy, especially within scenarios involving interpersonal communication. We conducted two exploratory co-design sessions with six older adults to investigate the design of rule-based access control systems for social robots, involving tasks such as 'narrative prototyping.' We analyzed participants' artefacts and discussions from these sessions to identify themes pertaining to: (1) their desire to be active users of robot privacy mechanisms; (2) challenges and opportunities for 'translating' privacy desires into rule-based robot policies; and (3) normative concerns regarding the integration of a robot within the shifting and interpersonal privacy dynamics of old age. We outline high-level design implications following from our analysis; our findings contribute to a growing, relational understanding of what it means to design privacy-conscious social robots for older adults.

GI'26, Waterloo, ON, Canada

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
https://doi.org/XXXXXXX.XXXXXXX

CCS Concepts

• **Human-centered computing → Participatory design**; **Empirical studies in interaction design**; • **Security and privacy → Social aspects of security and privacy**; **Privacy protections**.

Keywords

social robots, older adults, privacy, access control, co-design

ACM Reference Format:
Clara Y. Xi and Lora Oehlberg. 2026. Rules for Robots: Co-Designing Privacy-Conscious Social Robots with Older Adults. In *Proceedings of Graphics Interface 2026 (GI'26)*. ACM, New York, NY, USA, 16 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Human-robot interaction researchers have spent decades investigating how social robots could assist or enhance the lives of older adults (defined by Statistics Canada as those aged 65 years or older [72]). Social robots could be used to support the delivery of healthcare, monitor for health or safety emergencies, assist with household tasks, give reminders, and even provide companionship to their older adult users [12, 18]. Despite its long history, the gerontechnological project of social robot development shows little sign of slowing down. Contemporary social robots, such as Intuition Robotics' ElliQ [29] and Ageless Innovation's Joy for All line of robotic animals [3] have already seen large-scale deployment by governmental assistance agencies in the United States [17, 71], and are readily available for purchase on the consumer market.

However, even as social robotic development moves forward, issues surrounding privacy persist as key design problems and ethical challenges. Responsive companion robots such as ElliQ are always on and always 'listening,' storing vast amounts of private information about their users in their homes [13, 68]. The type of this private information ranges greatly, from health data to intimate moments and confidences, implicating user privacy along multiple dimensions [15, 20]. As physically embodied, autonomous social actors situated within people's personal spaces, social robots introduce distinct challenges beyond those already identified in familiar contexts such as smart home or Internet-of-things technology. The risks social robots present extend beyond pervasive data collection as social interaction abilities raise new and notable concerns. Consider a scenario where a person approaches a social robot to ask for some information about that robot's user (see Figure 1)—how should the robot determine whether or not to grant that person access to that information in that specific moment? What mechanism should govern a robot's disclosures and allow it to recognize and prevent a privacy violation from occurring, without depriving it completely of its ability to communicate?

Considering older adult users adds further perspective and urgency to these concerns. Researchers have repeatedly shown that privacy is a major concern among older adults who would otherwise be interested in social robots [69]. Older adults also have unique privacy management concerns; they often must continually renegotiate interpersonal privacy dynamics and boundaries as they navigate new age-related vulnerabilities and needs for care [74]. Elder abuse, which includes psychological and financial abuse, is also pervasive, with as much as ten percent of older adults experiencing abuse in Canada [58]. Finally, issues surrounding autonomy and paternalism are keenly felt; the ability to exercise control is important to older adults, but this is not always reflected in technologies designed for them [8].

Our work aims to address various topics within this web of concerns. From August to December 2024, we conducted a series of co-design sessions with six older adult participants, exploring issues at the intersection of social robotics and personal privacy management. As part of this series, we included two sessions which were focussed on developing concrete, specific design ideas for an access control system for a social robot. Our choice of a co-design approach is deliberate, aligned with historical calls for the increased, meaningful engagement of older adults across all stages of gerontechnological social robot design [23].

Our work in this paper contributes towards a greater understanding of the design challenges and opportunities for older adults' privacy management with a social robot. Through a thematic analysis of our participants' design artefacts and discussions from the two key co-design sessions, we identify three overarching themes pertaining to: our participants' desire to be actively engaged users of robot privacy mechanisms; their identification of challenges and opportunities for 'translating' real-world privacy expectations onto a robot; and their speculations regarding the integration of social robots within the shifting and interpersonal privacy dynamics of old age, particularly with regards to issues of increasing dependence and relational control. Our discussion of our themes, interwoven with specific insights and design contributions from our participants, serves as a valuable starting point for future inquiry towards the meaning and design of privacy-conscious social robots for older adults.

2 Background

2.1 Privacy Frameworks

Our work is theoretically informed by two highly influential privacy frameworks: Communication Privacy Management (CPM) theory [54, 56] and Contextual Integrity (CI) [49, 50].

2.1.1 *Communication Privacy Management theory.*

CPM aims to provide a comprehensive framework for understanding how people regulate the disclosure of private information in interpersonal contexts. It uses a boundary metaphor to illustrate how private information becomes more collectively managed once it is shared. Information that is not shared with anyone stays within that individual's privacy boundary. However, as soon as it is shared, it moves into a collective, co-owned privacy boundary. When more people are involved in a privacy boundary, collective management of the information within the boundary becomes more difficult as it demands negotiation and agreement between a larger number of co-owners.

CPM is further based around the notion of *privacy rules* as the mechanism by which people expect privacy boundaries to behave. These rules can be implicitly or explicitly understood, and be influenced by a number of factors such as cultural norms or situational considerations [55]. Privacy rules govern a variety of conditions surrounding boundary management, specifying the exact conditions in which flows of information in and out of the boundary are appropriate. CPM also gives explicit consideration to how systems of social relations readjust privacy rules in response to disruptions when individuals are dissatisfied with how information has been flowing. Boundary recalibration is an interactional process, requiring active communication between individuals to negotiate shared expectations and prevent future turbulences.

Despite originating as a theory of human-human communication, CPM has been extended and verified in the domain of human interactions with conversational agents such as chatbots [40, 67] and smart speakers [34]. These previous empirical findings suggest that people do view non-human conversational agents as co-owners of their private information, with the associated expectations for privacy management that CPM implies.

2.1.2 *Contextual Integrity.*

Nissenbaum's CI framework was developed "as an attempt to understand what people saw threatened by novel sociotechnical practices" [49]. It conceives of privacy as informational flows, and maintains that privacy is preserved when such flows adhere to collectively held *privacy norms* (analogous to *rules*). Malkin [46] elaborated on the concept of the privacy norm by illuminating the fact that "there may not be single true norm for a given situation," and that norms themselves are subject to rapid evolution and change: in particular those relating to novel information-collecting technologies. According to Malkin, "norms cannot be assumed or derived from first principles but must rather be gleaned from the real world"; that is, the only reliable way to identify privacy norms is through examining people's real-world attitudes, beliefs, and expectations about privacy.

According to CI, privacy norms for a given flow are determined by five contextual factors:

- *subject* (who the information is about);
- *sender* (who or what is sharing the information);
- *recipient* (who is receiving the information);
- *information type* (what sort of information is being shared); and
- *transmission principle* (the stipulations imposed on the information flow).

The first four parameters are relatively self-explanatory. The *transmission principle* parameter encompasses "the conditions or constraints that restrict information flow or limit it to specific circumstances" [46]. An example of a transmission principle is the subject's consent. It is worth noting that transmission principles are "structurally generative" [49]; for instance, Apthorpe et al. [7] generated and identified a number of transmission principles related to information flows in smart home technology, such as "if the information is used to perform maintenance on the device," "if the information is used to provide a price discount," and "if the information is stored indefinitely." In CI, each of the five factors must be specified in order to make a determination of the privacy norm at play for a given information flow.

2.2 Designing Private Social Robots

Currently, much existing work exploring the design of privacy-conscious social robots adopts an informational view of privacy, which Lutz et al. [43] define as encompassing "social robots' data collection and processing capacities." Grabler and Koeszegi [26] note that informational privacy is predominantly emphasized in current research relating to the design of privacy-sensitive social robots. Examples from the literature include technical strategies for the inhibition of data collection, such as through sensor adjustment (e.g., [16, 19, 35]). Prevention of information collection is a common theme in research involving privacy-preserving conversational agents in general, which includes other technologies such as chatbots and voice-activated personal agents. Example strategies include wake words and recording indicators [57] as well as the preprocessing of user prompts to contain less private information [48]. However, these efforts offer little in the way of guidance for maintaining privacy *after* sensitive information has already been collected, and say even less about appropriately handling private information in interpersonal, multi-user contexts. The privacy tensions of multiple users within a shared private space have been explored by Lau et al. [36] and Luria et al. [42] in the context of voice-activated personal agents. We build on their work through a deeper exploration into the interpersonal contexts of concern for older adults, highlighting the position of older adult users as active agents in implementing their own privacy preferences.

So far, we find that Tang et al.'s work on developing a rule-based privacy system for a social robot serves as most significant effort towards the design of an interpersonally private social robot to date [73]. Having been informed by CPM and CI, the authors describe their design of a robot "privacy controller" which extracts tuples of contextual metadata from overheard disclosures and runs them through an internal set of "privacy rules" to assign a privacy level ("low," "medium," or "high") to the disclosures, which then governs the robot's decision-making on whether or not to reveal the information to others. Tang et al.'s work, however, is not specifically oriented to the operational setting of a social robot used to support an older adult. Moreover, their privacy controller prototype relies on a fixed set of privacy rules (generated through crowdsourcing) that cannot be modified by a user, leaving the question of how users could interact with these rules unaddressed. In contrast, our work is specifically interested in exploring older adults as primary users of social robots and in asking the question of how they might wish to directly interact with their robot's privacy system.

2.3 Co-Designing Technologies With Older Adults

Co-design refers to a design approach where end-users are engaged as active participants and co-creators within a design process [65]. It is guided by an acute awareness of power dynamics between researchers/designers and end-users, recognizing it as a key issue within technology design [64]. Multiple commentators have historically remarked on this power dynamic as applied to social robot designers and older adults [70]. In a 2014 literature review, Frennert and Östlund found that most articles "described older people as objects and did not involve them as subjects in their research," and that social robot designers were often influenced by negative stereotypes of older adults, leading to an outcome where older adults' "matters of concern are not identified in the design process" [23]. They argued for the increased use of co-design, stating that through co-design, "stereotypical views of robots can be undermined and a clearer understanding of what social robots of today really can and cannot do can be achieved."

There exists a small but growing body of work in the human-robot interaction literature surrounding the use of co-design with older adults [60]. The existing work emphasizes the success of the co-design approach when paired with appropriate method design and researcher support. The use of diverse design activities allows older adult participants to think and design through multiple modalities [51], while providing opportunities for iteration allows participants to refine older ideas and build a sense of continuity through the study [38, 51]. Finally, proper scaffolding is needed to enhance participants' ability to exercise designerly skills [28] and support the core co-design process of mutual learning [37]. To date, however, the challenge of explicitly addressing the design of a privacy-conscious social robot has yet to be tackled through co-design. While older adults occasionally brought up issues of privacy in previous co-design studies (such as in [51]), privacy has never been the focal point of investigation.

When we turn to the literature on the co-design of usable privacy and security mechanisms, we find that this body is even more limited, particularly with regard to older adult end-users. Current co-design research tends to be centred around end-users who are young adults [62, 76] or children and adolescents [4, 39, 53]. There exists one notable co-design study performed by Yao et al. [75] which engaged a subgroup of older adult participants in the investigation of privacy mechanisms for smart home technologies. However, the authors' reporting of their results addresses all end-users of smart home technologies in general, subsuming discussion of whatever unique design implications that may have existed for the older adult user group.

3 Methodology

As part of our series of co-design sessions which took place from August to December 2024, we conducted two co-design sessions specifically with the aim of developing design ideas for an access control-based privacy system for a social robot. These two sessions were the final two sessions of the series, and thus relied on the participants' attunement to the privacy issues of social robots as well as the development of their designerly skills established through earlier sessions. Figure 2 provides a visual summary of the five sessions from our overarching co-design project, which involved multiple activities including interviews, robot demonstrations, vignette reading and writing, and collage-making. Our co-design sessions were approved by the University of Calgary Conjoint Faculties Research Ethics Board (REB-0530). We provide a quick overview of the first three co-design sessions to provide methodological transparency and situate our participants' experiences before describing the methods of our final two sessions, which form the basis of this paper's analytic contributions, in detail.

We recruited our participants through a posting to a local seniors' Facebook group and an in-person recruitment event during a seniors' gathering at a community centre. All of the participants took part in all of the co-design sessions, which took place in their homes. Each session except for the first took roughly two hours to complete; the first session took roughly one hour. Sessions with participants took place at least one week apart from another. Participants were compensated with a $25 grocery store giftcard at the end of every session.

3.1 Prior Sessions

The first session with each of our participants consisted of an initial semi-structured interview. During the interview, we asked participants general questions about their daily lives and routines, as well as questions about their familiarity with robots and other forms of technology. The initial interview allowed us to collect basic information about our participants and establish initial rapport.

The second session was centred around introducing social robots to our participants. We began the session with a series of live and video robot demonstrations, where we presented participants with a physical NAO robot [25] followed by videos (sourced from YouTube) of Aldebaran Robotics/Softbank Robotics's Pepper [5, 24, 52], Fraunhofer IPA's Care-O-Bot 3 [22, 27], Ageless Innovation's Joy for All cat [2, 3], and Intuition Robotics's ElliQ [30, 29]. These robots were chosen for their variety in form and function to showcase the diversity of social robots. After discussing the robots from the demonstrations with our participants, we launched into a co-design task where we presented participants with a worksheet to write down design ideas for their own 'ideal' social robot, addressing broadly desired functionality and presence within daily life. We designed this task to give participants the opportunity to reflect on their new knowledge of social robots and be eased into the creative, generative aspect of co-design.

We introduced the intersection of privacy issues with social robotics in our third session. The first task of this session involved reading and writing short privacy vignettes. We began by presenting participants with two vignettes that the we prepared featuring a social robot potentially violating its user's privacy. We then asked participants to write their own vignettes depicting a hypothetical privacy scenario of their own choosing. For the second task of this session, we asked participants to create collages to visually and tactilely explore the emotional implications of a social robot's inappropriate privacy management.

3.2 Thinking About Access Control

Situated within our wider co-design project, the focus of our final two co-design sessions was to converge [41] participants' thinking towards the specific design goal of an access control system for a social robot. Thus, the aim of the fourth and penultimate session was to build our participants' understanding of access control and begin to generate concrete, specific design ideas for an access control system for a social robot.

Access Control Interactive Presentation. We spent the first half of the session introducing participants to the basic concepts of access control as well as the privacy frameworks of CPM and CI through an interactive presentation. The roughly forty-minute presentation was delivered using a slide deck that was shown to participants on a laptop. To keep participants engaged and support their active learning of the concepts, we made the presentation interactive through a number of signposted discussion questions to help participants reflect on or exercise their understanding of the material. For example, we asked participants to list potential privacy rules they followed in their own lives after introducing this concept from CPM; additionally, we asked them to practice identifying CI factors in the privacy vignettes that both the first author and the participants wrote in the third session. Highly technical ideas and jargon were avoided as much as possible since our main goal was to convey a general conception of the management of private information as governed through *contextually-informed privacy rules* that could be 'translated' unto a robot in the form of access control rules or policies.

Access Control Design Activity. In the second half of the session, we asked participants to generate ideas for an access control system for a social robot geared towards interpersonal privacy management. To scaffold our participants' thinking, we designed a worksheet (see Figure 3) that broke up a hypothetical robot privacy system into various aspects. These included input modality, the ability to handle ambiguous or conflicting rules, output during privacy-sensitive interactions, and considerations of the rule structure proper. The worksheet maintained the underpinning that CPM-informed privacy rules were analogous to a robot's 'privacy' or access control rules, and reinforced core concepts of CI by reference to the 'factors' by which rules should be constituted. After participants filled out the worksheet, we discussed their responses with them.

3.3 Consolidating Ideas and Co-Design Debrief

Our fifth and final co-design session was focussed on giving participants the chance to iterate or reflect on their design ideas from the previous session. Furthermore, as part of the final session, we sought to facilitate the consolidation of our participants' design

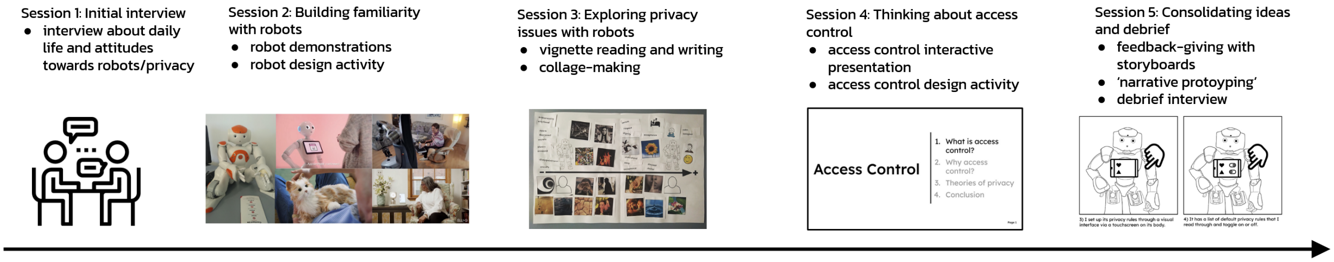

Figure 2: Visual summary of our full co-design project. Sessions 4 and 5 (excluding the debrief interview) are the main focus of this paper.

ideas and experiences from the previous session and beyond into a single, contextually-situated prototype.

Feedback-Giving With Storyboards. After each instance of the fourth session concluded, the first author would create three to four storyboards based on design ideas participants developed during the access control design activity. The first author handpicked ideas to form storyboards around based on their own subjective interpretation of which ideas they found particularly salient or believed could be clarified further following concretization into a storyboard format. The purpose of the storyboards was to probe participants into iterating or clarifying their ideas from the previous session. We showed participants the storyboards we had prepared for them at the beginning of the final session, spending time to discuss each one.

'Narrative Prototyping'. What we refer to as the 'narrative prototyping' activity made up the final design activity of the study. We came to the idea of the use of narrative or storytelling as a prototyping method as there were a few things about a privacy management system we wished to explore in a rich, temporal context, such as how a user would manage privacy rules on the robot and how the robot would behave in privacy scenarios. To scaffold the creation of a narrative, we created a worksheet for this design activity based off of a fairy tale story spine format. We adapted Kenn Adams's story spine, originally used as a tool for improv theatre [1], due to its ability to facilitate the rapid construction of conventionally-structured narratives. Adams's story spine has been previously used in participatory research to collect narratives from a diverse set of stakeholders (e.g., [61]).

The narrative prototyping activity was split into two parts. The first part of the activity was designed to accustom participants to the story spine format, where we presented participants a privacy scenario adapted from one of the preprepared vignettes from the third session, modified to fit the story spine. Participants were asked to fill in a blank prompt to complete the narrative. For the second part of the activity, participants were given blank story spine headers accompanied by more specific prompts, such as: "What events happen that could lead to the robot causing a privacy violation?", "What privacy rule needs to be given to the robot?", and "How is the privacy violation averted?" (see Figure 4).

Pseudonym (self-chosen)	Age (years)	Gender
Karen	75	F
Bill	75	M
Valerie	74	F
John	82	M
Don	71	M
Historic	82	M

Table 1: Participant information.

3.4 Participants

A total of six participants took part in our project; a summary of the participants may be found in Table 1. Of the six participants, two couples—Karen and Bill, and Valerie and John—took part in the co-design sessions as dyads. The other two participants, Don and Historic, participated in the sessions individually. None of our participants had prior experience with social robots; they overall had very little knowledge of social robots before the second co-design session. They were generally comfortable with technologies such as smartphones and tablets. Although a minority were affected in their day-to-day lives by physical impairments, our participants generally viewed themselves as healthy and active members of their communities.

3.5 Data Analysis

All sessions were audio recorded with participant consent. We used reflexive thematic analysis [11] to analyze the combined corpus of the verbal discussions which took place during the final two sessions as well as the participant-generated design artefacts produced by those sessions (namely, the text-based access control design worksheets and the narrative prototypes). We transcribed our audio recordings as well as the text of the design artefacts by hand, and performed our analysis on these transcripts. We followed the steps outlined by Braun and Clarke in [10] to complete our analysis. First, we familiarized ourselves with the data by reading it multiple times. Then we performed complete coding on the data. This was done using the open-source qualitative coding software QualCoder [59]. Then we examined the coded, collated data and searched for organizing thematic patterns. In developing our themes, we returned

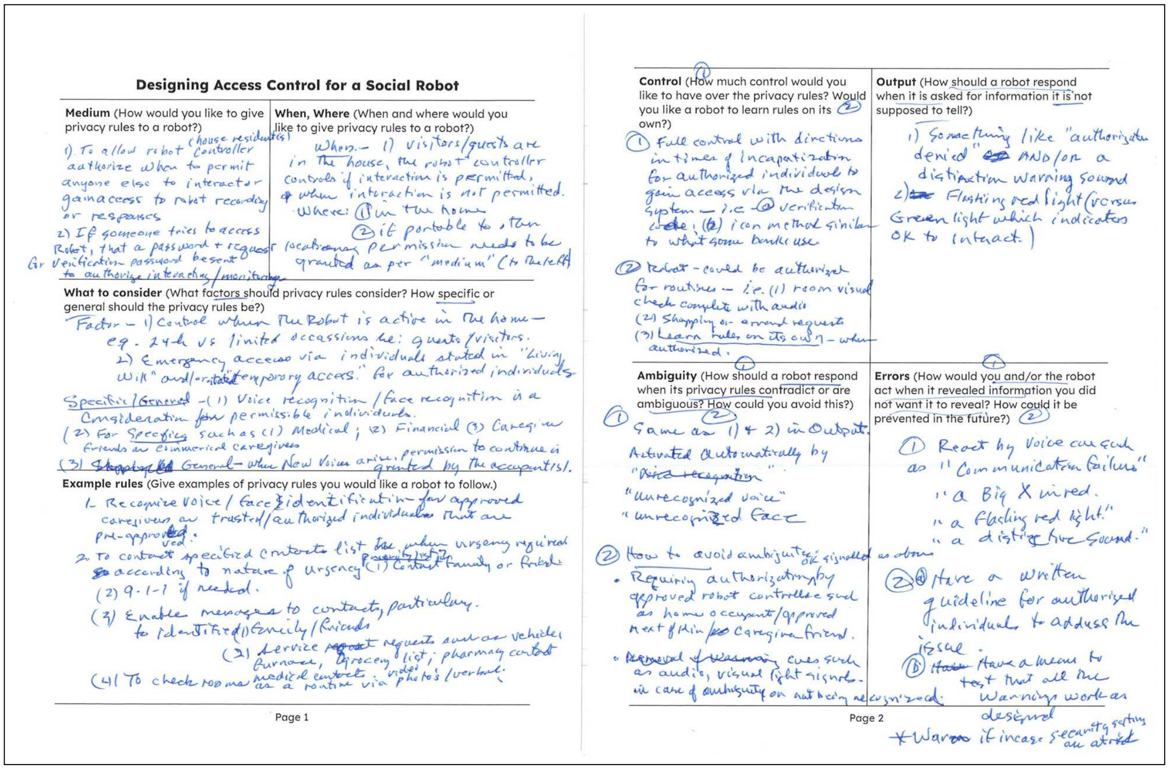

Designing Access Control for a Social Robot

Medium (How would you like to give privacy rules to a robot?)
1) To allow robot controller (house residents) authorize when to permit anyone else to interact or gain access to robot recording or responses
2) If someone tries to access Robot, that a password + request for verification password be sent to authorize interacting/monitoring.

When, Where (When and where would you like to give privacy rules to a robot?)
When. – 1) visitors/guests are in the house, the robot controller controls if interaction is permitted, when interaction is not permitted.
Where: 1) in the home
2) if portable to other locations permission needs to be granted as per "medium" (to Robot)

What to consider (What factors should privacy rules consider? How specific or general should the privacy rules be?)
Factor – 1) Control when the Robot is active in the home – eg. 24-h vs limited occassions ie: guests/visitors.
2) Emergency access via individuals stated in "Living Will" and/or temporary access. for authorized individuals
Specific/General – (1) Voice recognition/face recognition is a consideration for permissible individuals.
(2) For Specifics such as (1) Medical; (2) Financial (3) Caregiver friends or commercial caregivers
(3) General – when New Voices arise, permission to continue is granted by the occupant(s).

Example rules (Give examples of privacy rules you would like a robot to follow.)
1. Recognize voice/face identification for approved caregivers or trusted/authorized individuals that are pre-approved.
2. To contact specified contacts list when urgency required according to nature of urgency (1) Contact family or friends
(2) 9-1-1 if needed.
(3) Enable messages to contacts, particularly to identified family/friends
(3) service requests such as vehicles, furnace, grocery list, pharmacy contact medical contacts: video
(4) To check rooms as a routine via photo's/verbal.

Page 1

Control (How much control would you like to have over the privacy rules? Would you like a robot to learn rules on its own?)
1) Full control with directions in times of incapacitation for authorized individuals to gain access via the design system – i.e. @ verification code; (b) icon method similar to what some banks use.
2) Robot – could be authorized for routines – i.e. (1) room visual check complete with audio
(2) Shopping or errand requests
(3) Learn rules on its own – when authorized.

Output (How should a robot respond when it is asked for information it is not supposed to tell?)
1) Something like "authorization denied" AND/OR a distinctive warning sound
2) Flashing red light (versus Green light which indicates OK to interact.)

Ambiguity (How should a robot respond when its privacy rules contradict or are ambiguous? How could you avoid this?)
1) Same as 1) + 2) in Output. Activated automatically by "unrecognized voice" "unrecognized face"
2) How to avoid ambiguity or signalled as above
• Requiring authorization by approved robot controller such as home occupant/approved next of kin or caregiver friend.
• Approval of warning cues such as audio, visual light signals in case of ambiguity or not being recognized.

Errors (How would you and/or the robot act when it revealed information you did not want it to reveal? How could it be prevented in the future?)
1) React by voice cue such as "Communication failure" "a Big X in red." "a flashing red light." "a distinctive sound."
2) a) Have a written guideline for authorized individuals to address the issue.
b) Have a means to test that all the warnings work as designed
* Warn if increase security setting are altered

Page 2

Figure 3: An example of a filled-out access control design worksheet.

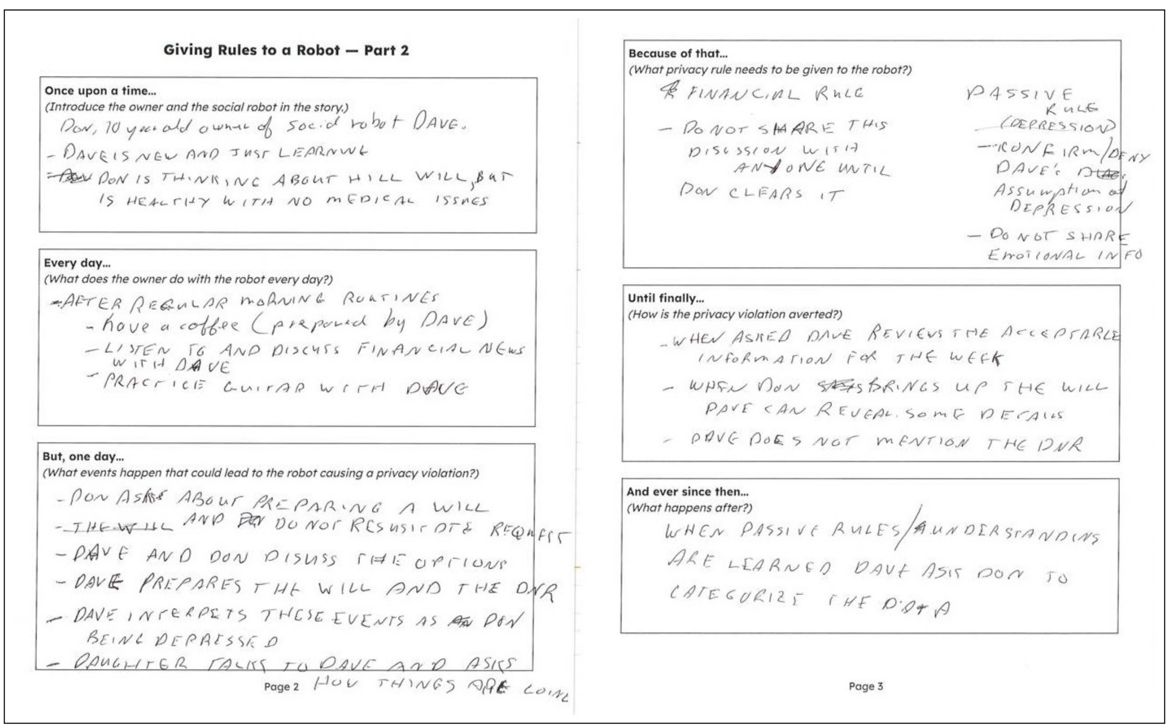

Giving Rules to a Robot — Part 2

Once upon a time...
(Introduce the owner and the social robot in the story.)
DON, 70 year old owner of social robot DAVE.
- DAVE IS NEW AND JUST LEARNING
- DON IS THINKING ABOUT HIS WILL, BUT IS HEALTHY WITH NO MEDICAL ISSUES

Every day...
(What does the owner do with the robot every day?)
- AFTER REGULAR MORNING ROUTINES
- have a coffee (prepared by DAVE)
- LISTEN TO AND DISCUSS FINANCIAL NEWS WITH DAVE
- PRACTICE GUITAR WITH DAVE

But, one day...
(What events happen that could lead to the robot causing a privacy violation?)
- DON ASKS ABOUT PREPARING A WILL
- AND DO NOT RESUSICITATE REQUEST
- DAVE AND DON DISCUSS THE OPTIONS
- DAVE PREPARES THE WILL AND THE DNR
- DAVE INTERPRETS THESE EVENTS AS DON BEING DEPRESSED
- DAUGHTER TALKS TO DAVE AND ASKS HOW THINGS ARE GOING

Page 2

Because of that...
(What privacy rule needs to be given to the robot?)
FINANCIAL RULE
- DO NOT SHARE THIS DISCUSSION WITH ANYONE UNTIL DON CLEARS IT
PASSIVE RULE (DEPRESSION)
- CONFIRM/DENY DAVE'S ASSUMPTION OF DEPRESSION
- DO NOT SHARE EMOTIONAL INFO

Until finally...
(How is the privacy violation averted?)
- WHEN ASKED DAVE REVIEWS THE ACCEPTABLE INFORMATION FOR THE WEEK
- WHEN DON BRINGS UP THE WILL DAVE CAN REVEAL SOME DETAILS
- DAVE DOES NOT MENTION THE DNR

And ever since then...
(What happens after?)
WHEN PASSIVE RULES/UNDERSTANDINGS ARE LEARNED DAVE ASKS DON TO CATEGORIZE THE DATA

Page 3

Figure 4: An example of a filled-out narrative prototype.

to our data multiple times, re-performing the process of familiarization, coding, and pattern analysis until we reached our final themes.

3.6 Positionality Statement

The analytic lens and initial groundings of this work are primarily driven by the positionality of the first author. I (the first author) developed and ran the co-design sessions as a MSc student only a year after I had moved into the city in which the sessions took place, where I had no preexisting social network. I developed the research directions of the overall co-design project by combining my academic interest in social robotics for older adults with my prior experience in data privacy research. I facilitated each session alone, and my interpretive stance is deeply shaped by my sense of rapport with the project's participants. I additionally held a voluntary, informal gathering with the project's participants in June 2025 where I debriefed them on this work's early analytical findings. As a researcher, I am greatly motivated to continue to work, learn, and foster collaborative relationships with older adults and their communities.

4 Results

Through their access control design worksheets, their narrative prototypes, and their verbal discussions of these artefacts along with the storyboard probes, our participants provided a rich array of design suggestions and speculative insights. We identified three themes pertaining to how our participants envisioned privacy management in social robots. These three themes were: *(1) Active Engagement with Privacy Systems; (2) Challenges and Opportunities for Translating Privacy Desires;* and *(3) Supporting Shifting and Social Dynamics.*

A note on formatting participant quotations: wherever an extract from an artefact (i.e., the access control design worksheet or the narrative prototype) is quoted, it is enclosed in guillemets («»). This is to distinguish such extracts from quotes drawn from verbal discussions, which are enclosed in double quotes (“”).

4.1 Active Engagement with Privacy Systems

Participants were highly concerned with how users could understand and validate the state of a social robot's privacy system, as well as exercise as high degree of control over it. We report on this theme through two subthemes: *Understanding and Verifying the Privacy System*; and *User Control and Awareness of Robot Processes.*

4.1.1 Understanding and Verifying the Privacy System. Participants wanted users to be able to freely review the access control or privacy policy in place on their robot in order to build up their understanding of the policy. Such a review process could be enabled through the use of a visual representation of the robot's privacy policy. In the access control design activity, Bill suggested that *«the robot should have a built-in menu of privacy rules that you can accept or turn off»* (see Figure 5). Meanwhile, Don independently discussed the need for a visual representation of a robot's privacy rules via a *“grid that says well here's my hierarchy of information and here's all the people that can access it,”* calling to mind the primitives of an access control list or matrix [66]. In a similar manner to Bill's ‘menu,’ Don also imagined that users should be able to use the

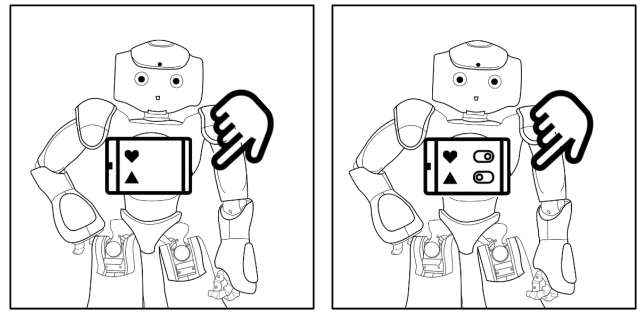

Figure 5: Fragment of a storyboard prepared for Bill demonstrating his desire for an explicit, easily accessible means to view and adjust privacy rules.

visual representation to review rules at an atomic level. For Don, reviewing the rules through an easy-to-access visual representation was important as it *“would let me understand where I stand security-wise.”* Far from viewing themselves as passive beneficiaries, our participants demonstrated their keen interest in taking active, ongoing command over their own privacy.

Participants also brought up how users should be given the means to audit their robot's privacy behaviour. Auditing could take the form of short, verbal indications of disclosure, as Valerie had suggested (*«the robot would tell me he shared information»*), to an extensively documented *“audit trail of information that's been revealed, what information's been revealed to who, when,”* according to Don. Don positioned auditing as a necessary compliment to access control policies and mechanisms in order to attain guaranteed security: *“the ACL [access control list] shows you what is supposed to happen and the audit trail tells you what did happen.”* Furthermore, for Don, auditing the robot's interactions was not only a way to validate its access control but to also reinforce a sense of control over the robot: *“I don't think you can turn your life over to the robot, [not knowing] what's going on.”*

Verification and trust in a robot's privacy system was a particularly salient issue for John. John did not believe it was possible to perform *a priori* verification on a robot's privacy rules (*“I don't think it's humanly possible”*), stating: *“it's quite amazing still the number of things you cannot anticipate until you actually test them and the failure brings you to the fact that we omitted this.”* He expressed worry for a situation in which a flaw in the privacy rules may go unnoticed for *“a long period of time [...] could be many years [until] suddenly there was some quirky situation that revealed this was uncovered and led to some embarrassing outcome or some seriously risky one.”* John's comments reflect a fundamental challenge known to the theoretical access control literature, where *a priori* formal verification of an access control model is understood to be generally intractable [32].

However, other participants were more hopeful about the ability of older adults to put their trust into a social robot's privacy system. In Bill's narrative prototype, witnessing his social robot act in accordance with the privacy rules he gave it during a critical interaction served to validate its access control system and build its user's trust in it:

> Bill: *"The rule worked and observing that the rule worked made us think, ah, we should probably add a few other specific rules."*
> Clara (first author): *"Kind of developing that trust in the robot that it's going to keep doing what you asked it to?"*
> Bill: *"Yeah, and also trust in the way the rule was formulated."*
> Karen: *"That the robot understood."*
> Bill: *"Yeah. That the formulation was correct leading to the proper response by the robot. It was like this test [...]"*

4.1.2 User Control and Awareness of Robot Processes. Participants stressed their expectations for a high degree of user control over a social robot's privacy system. In the access control design activity, We asked participants about the level of control they would like to have over a robot's privacy rules—unanimously, participants answered the prompt with 'total,' 'full,' or 'complete' control.

The issue of a robot autonomously 'learning' or inferring new privacy rules without its user input, which we had specifically prompted participants on during the access control design activity, had mixed reactions. Some participants, such as Don and Historic, were open to the idea of a robot autonomously developing privacy rules on its own. Don conceived of the learning process as *"observing [my] actions and interactions with other people and learning from that as opposed to me explicitly saying this is something you should not do [...] sort of an AI kind of approach to it."* In response to the storyboard we prepared featuring this learning process (see Figure 6), Don commented on its utility in forming a comprehensive set of privacy rules: *"I think that's important because I can't think of everything. So it's good that the robot notices things and tries to learn."*

Other participants were more hesitant or distrustful towards the idea of a social robot developing privacy rules for the user on its own. In response to our prompt about a robot learning rules on its own, Valerie simply stated, *"no, I'd like to be there with him or her."* Karen and Bill also expressed apprehension at the opaqueness of an autonomous rule-learning process, discussing how it could lead a robot to infer incorrect privacy requirements:

> Bill: *"The robot would overhear you chatting with your girlfriend and go, 'Ah, it's okay to share this information, because she just shared this information.'"*
> Karen: *"Meaning the information is okay to share but [the robot might] not necessarily understand that it's in the context of certain friends."*
> Bill: *"Right, yeah."*
> Karen: *"This is a little, uh—"*
> Bill: *"Worrisome."*
> Karen: *"Yeah."*

Overall, participants kept returning to the importance of keeping a user aware of what their robot has learned or inferred about them. Karen and Bill concluded that if a robot were allowed to learn privacy rules on its own, it would have to keep its user in the loop and acquire their permission before implementing new rules:

> Bill: *"So if the robot formed new rules, those would have to be—"*
> Karen: *"On the screen."*
> Bill: *"Communicated to the owner [...] The robot would have to tell you what it thinks the rule is so you can go, 'that's not right.'"*

Likewise, Don and Historic—the two participants who were initially amenable to the idea of a robot learning rules on its own—were also quick to highlight the need to acquire user authorization before implementing new rules. Don further explained that while it was good for a robot to *"proactively [say] 'I have a new rule and want to examine it,'"* he as a user would still want to be in the seat of control in the final analysis: *"there's the other way also, that I would observe the robot for a few days and I would wanna say, 'Okay, hold it, there's a few things where it didn't work,' so I wanna be the one to be proactive, not the robot."*

For John and Don, the need to loop in the user into their robot's learning processes was extended to include all information that a robot had acquired or inferred about its user. John suggested that this could happen as an in-advance negotiation intertwined with policy formation, using an example of sharing information to facilitate *"intergenerational bonding"* between an older adult and their grandchild:

> *"The robot makes a rule, 'Why don't you let me get these details [about your cooking] and I could tell a story to your granddaughter?' That's engaging the grandmother in this where the grandmother might be very uncomfortable with any technological intrusion in her secret recipes [laughs]."*

On the other hand, Don imagined that a user could review what their robot has learned or inferred about them in a regularly-scheduled (e.g., weekly) review process with the robot, and categorize new items of information and develop the robot's privacy rules. This was illustrated most clearly in his final narrative prototype, where the robot 'Dave' asks Don to confirm, categorize, and form a new rule around what it has inferred about Don's emotional state (*«when passive rules/understandings are learned Dave asks Don to categorize the data»*).

4.2 Challenges and Opportunities for Translating Privacy Desires

Participants recognized the inherent challenge of translating a user's privacy expectations into the form of an explicit privacy policy a social robot could follow. We discuss how participants addressed this through two subthemes: *Policy Expressivity;* and *Scaffolding Policy-Oriented Thinking.*

4.2.1 Policy Expressivity. When we asked participants to come up with example policy rules they would like a social robot to follow, the rules they provided tended to be general and coarsely-grained, as opposed to dealing with specific people or items of information. Examples of generally-worded rules include: *«restrict the robot's use to its owner only»* (John) and *«do not share what I do, where I go, or what I say»* (Valerie). These sample rules given by John and Valerie

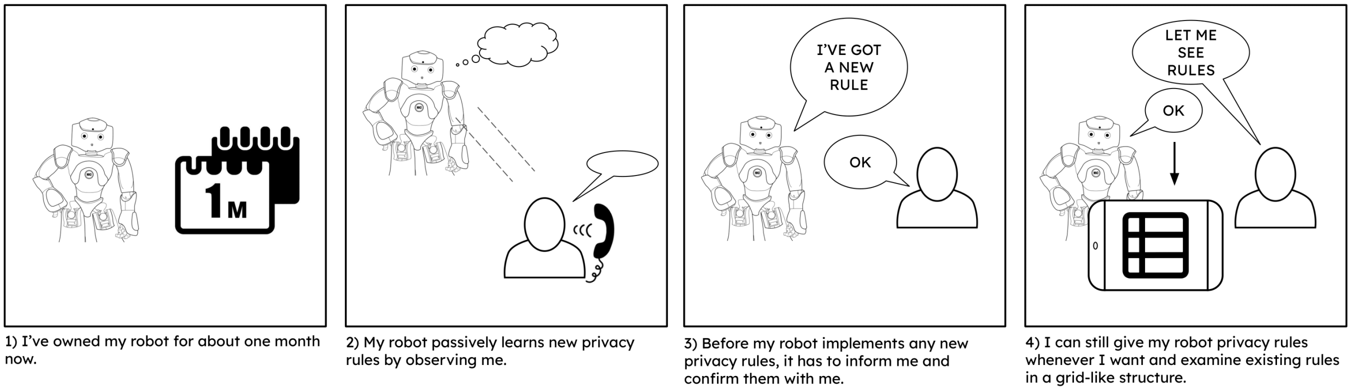

Figure 6: Storyboard prepared for Don demonstrating his desire for a robot to 'learn' rules through observation while still maintaining full user control and oversight.

are especially austere, dovetailing with comments made by John, Bill, and Don regarding their desire to approach the handling of their private information in a conservative, 'better-safe-than-sorry' fashion (*"the default is don't reveal"* [Don]). Don conceded that the eventual loosening of a robot's privacy policy could be developed over time in a manner similar to how privacy boundaries between people loosen as they develop trust between each other: *"it really could be really locked down to begin with, then you slowly open it up [...] same as you would with a person, really."*

Returning to the matter of coarsely-grained rules more generally, Don explained his motivation for the use of *"cover-all"* privacy rules as follows:

> *"'Do not share this discussion with anyone until I clear it' [...] that can happen quite often if you are talking to your robot. You don't want to make up a really defined rule, you want to say 'everything we just talked about is off-limits,' so it's sort of a cover-all rule and it's not on the fly."*

In this vein, coarsely-grained rules prevent insecurities in a privacy policy from appearing, as well as reduce the mental burden of rule management on the user by simply needing to express fewer rules to adequately protect one's privacy—a process which John otherwise foresaw as *"awfully laborious."*

When it came to expressing desired policy rules to a social robot, participants generally expressed these rules in highly natural language; nevertheless, these rules evidenced systematic consideration of the standard contextual factors derived from CI. In the investigative context of this study, the CI factors of *subject* and *sender* are already accounted for (with the main older adult user as the subject and their robot as the sender); this leaves the remaining three factors of *recipient, type of information,* and *transmission principle.* Participants near-unanimously incorporated recipient and type of information as an explicit factor in their policy rules. In particular, some participants, such as Bill, suggested how an organized system for rule creation could be based on structured matchings made between these two factors (see Figure 7):

> *«You should be able to specify individuals with whom information can be shared*
> *Information should be filed in categories and you can select what categories can be shared with whom»*

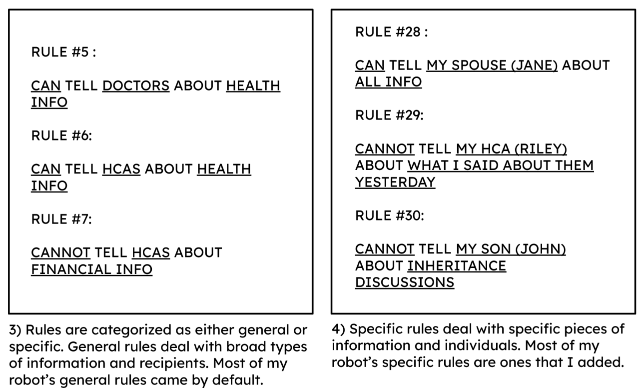

Figure 7: Fragment of a storyboard prepared for Bill demonstrating his proposed privacy rule structure based on the CI factors of recipient and type of information.

In contrast, while the CI factor of transmission principle was never mentioned explicitly as such, it was clear that participants also wanted a social robot to be able to take conditional factors on the disclosure into consideration. Some examples of transmission principles evident in participant suggestions included:

- approaching a topic through obfuscation or misdirection (e.g., *«only provide factual information»* [Don]; *«be vague in your answers»* [Valerie]);
- discussing a topic only after the user has broached it (*"none of this stuff I want the robot to reveal before I did"* [Don]);
- discussing a topic only when asked about it (*«provide info when asked—not proactively»* [Don]); and
- not revealing information learnt through hearsay (*"you could have a rule about hearsay where it could share hearsay or not"* [Bill]).

Most of these transmission principles were tied into issues of user awareness and control, with Don expressing concerns about *"a robot [. . .] that takes away all of your human contact with other people"* if strict limitations on the conditions of the robot's disclosures were not imposed. Transmission principles regarding obfuscation or misdirection may be necessary, as Valerie suggests, so that *"the person [interacting with the robot] isn't insulted or concerned that the robot won't talk to them."*

A transmission principle such as the one previously mentioned leaves room for interpretation—how vague is "vague" supposed to be? Dealing with natural language presents obvious technical challenges, which some participants commented on directly. Together, Karen and Bill discussed the challenge of making sure policy rules expressed in natural language pertaining to a certain information type would actually capture all information related to that type:

> Karen: *"If the robot is strictly going by exactly the word, for example 'inheritance,' they may not know that [the phrase] 'I'm leaving John this amount of money' is about inheritance [. . .]"*
> Bill: *"There's a lot of things that would be a part of that."*

On the other hand, Don discussed the challenge of making appropriate distinctions within a single information type: *"within each one of those information categories that we've defined here, there might be several [subcategories]."* He illustrated this with a concrete example based on the release of 'health information': *"If I just had a heart attack and I'm on the floor, yeah I don't mind if you call somebody [. . .] whereas if I went to the doctor and he said, 'Okay, you should be taking this medication,' that's a totally different thing, right?"*

4.2.2 Scaffolding Policy-Oriented Thinking. To help sensitize a new user to thinking about privacy rules explicitly, participants suggested various forms of structured interaction. A number of participants commented on the fact that it would be unlikely for a new older adult user to be immediately prepared to manage a privacy policy on their own. Don pointed out how he was able to talk about privacy controls fluently during the co-design study, but only because *"you and I have talked about it for four weeks and so my brain's a bit engaged in what it is, but a brand-new person who never thought about it, it's too much of a blank slate in some ways, so you have to have some structure."*

One key means through which this could be achieved involved the inclusion of 'default rules' that would be shipped with the robot from the get-go. As Karen explained, *"somebody who buys the new robot may not even think about privacy rules"*—but by having default rules evident upon setup, *"the person who buys it has to make those decisions and is made aware."* Bill commented that through careful consideration of what privacy rules should be included by default (perhaps through a similar co-design process), *"you would start to see what were common things that people would be concerned about and be able to develop specific rules so that there'd be less of a learning curve."* Similarly, Don also suggested the inclusion of 'standard' default privacy rules that should be 'obvious,' revolving around topics such as financial information.

To address the difficulty in developing a comprehensive yet complex privacy policy for a social robot, John described an idea for a levelling system that would scaffold the user's ability to input more intricate privacy policies:

> *"You might be able to load it up with a level one, level two, level three [. . .] It's in simplicity mode for the user's first three months let's say, and there may be some test program you can put in that's based on a scenario that is programmed to test the user's understanding [. . .]"*

Moving between these levels could be accompanied through targeted instruction: *"You go into level two very quickly and there's a roadblock that says, 'I think you should spend a week looking at an introduction to level two.'"* John's suggestions proffer a view of older adults as receptive to technological training, in contrast with stereotypically negative perceptions of older adults as technologically resistant and incapable.

In a different style of approach, Valerie was supportive of the the idea that a robot's privacy policy could be 'worked out' through verbal, conversational interaction, although she gave the caveat that *"it has to be a pretty intelligent robot."* Don was especially optimistic about the opportunities for guided policy creation through intelligent conversational interfaces, having witnessed advancements in them for himself:

> *"I was with my older sister [. . .] she spoke to ChatGPT as though she was talking to another person, just totally natural, and I was amazed. She was trying to find out the name of a lodge on a highway [. . .] ChatGPT said, 'Oh, it could be these,' and she said, 'No, no, it's not going in that direction' [. . .] [ChatGPT] says, 'Oh okay so it might be these three' [. . .] That's the kind of interaction where you sort of want to zero in on something [. . .] I'd want that sort of give and take and conversational style that would help me figure out what the rule is. It should be possible because you're doing it now [laughs]."*

4.3 Supporting Shifting and Social Dynamics

Aging is an inevitable process of change, and not always in ways we are necessarily comfortable with or able to face alone. Participants were keenly aware of the unique life circumstances faced by older adults, and readily discussed the potential place of social robots in responding to shifting circumstances and integrating into preexisting structures of social support. This theme is divided into two highly interconnected subthemes: *Changing Circumstances;* and *It Takes a Village.*

4.3.1 Changing Circumstances. The need for a social robot to adapt to shifting life circumstances and the new privacy norms that came along with them was a prominent theme in my participants' narrative prototypes. Participants did not shy away from portraying various changes or declines associated with aging, and readily discussed how the privacy behaviour and roles of a personal social robot should shift in response to these events.

In Karen's narrative prototype, the user, 'June,' suddenly collapses one day at home; the privacy policy on her robot, 'Annette,' prevents Annette from *«[revealing] June's activities or medications to anyone»*, information which *«emergency responders need [. . .] to properly respond to June»*. The user is sent to the hospital; once she returns, she loosens the privacy policy so that her robot is allowed to disclose medically-relevant information to certain individuals such as emergency responders and other healthcare personnel. Karen framed the events of her narrative as a wake-up call for the user to re-evaluate her privacy boundaries and realize her robot's potential instrumentality in supporting her health needs:

> *"June likes to be in control, so she didn't want anybody to know what medications she's on or if she's had an emergency or whatever. But June is smart enough to know that in the case of an emergency or medical issue it's better that Annette knows what drugs she's on and who to disclose it to. June's a fast learner."*

John's narrative prototype was also centred around a life change associated with old age—declines in cognitive ability. In John's narrative, the user 'Pat' is said to enjoy going on walks in the forest with his robot, 'Robbie,' taking pleasure in identifying different species of trees. One day, the robot inadvertently violates its user's privacy by revealing the progression of the user's memory loss in a conversation with one of his friends:

> *«On one recent occasion, Pat met an old friend that questions Robbie about the different roles and supports he provided for Pat. Robbie chose the example of Pat's remarkable knowledge of the many different tree species in the woods. Robbie proceeded to add that recently Pat had begun to misidentify or forget some of the different tree species.»*

John conceived of memory loss as an inevitability for older adults—*"typical [...] for all of us, whether we like it or not."* In order to prevent the *"embarrassing"* privacy violation depicted in his narrative, John suggested that the robot could perform some sort of *"self-test that would put an expiry date on [its] public behaviour."* The purpose of the test would be to check *"that [the user is] still consistent in their cognition";* in John's view, social robots who bear witness to their users' age-associated decline should be expected to adjust the privacy rules they follow over time in order to handle these topics with tact.

As the two narratives demonstrate, age-related life changes may result in either a tightening or loosening of the shared privacy boundary between the user and their robot, which may be dependent on the individual privacy preferences or needs of the user. These changes can be sudden or gradual in onset, and are often unwelcome as they are usually associated with a reduction of abilities and increased dependency on others. Highlighting this last point, Karen argued that, in turn, the amount of assistance social robots would be expected to provide would necessarily increase as well: *"I think robots either need to be able to offer more as the client declines or you'd have to get a whole different kind of robot."* A social robot designed for long-term use by older adults should be prepared to adapt nimbly and appropriately to the various life changes their users may encounter—not just in their privacy policy or behaviour, but in their holistic role in their users' lives.

4.3.2 It Takes a Village. Participants underscored the socially-situated nature of the robot and the user, noting the relational nature of privacy boundary management and negotiation. Karen elaborated on the events of her narrative prototype to explain how June might have gotten to the point of re-evaluating her privacy boundaries after having a conversation with her children:

> *"It would be appropriate if she has a decent relationship with her son, daughter, whatever, that they should discuss this with her and decide together, 'Okay, Mom, you don't like doing this but you've seen what happened in this emergency. So from this day forward we're going to need a little bit more information about what medical issues you have, what you're on, and even who [we] can share [the information] with in addition to the robot.'"*

Karen's explanation demonstrates the relational character of privacy boundary negotiation—rather than privacy rules being solely determined by the user and effected solely on their robot, they instead arise from the negotiated desires of the user and their circle of human relations, and are effected as a collective. This subtheme is closely intertwined with its partner subtheme *Changing Circumstances*, since life changes are what prompt boundary and relational renegotiations in the first place. As Karen's comment illustrates, it is unlikely that users and their relational circle will allow social robots to be a peripheral consideration during their interpersonal privacy management discussions. It seems more likely that social robots will be central players in these discussions, owing to their naturalized status as privacy co-owners and their closeness to their users themselves, especially if they will be tasked with keeping important information about them.

Privacy negotiations between an aging older adult and their relations can already be an uncomfortable and turbulent process [74]. It is even more unclear what this process might look like if a social robot becomes involved. Nevertheless, participants were hopeful that human adaptability and grace would allow the development of new privacy norms to take place even in an unfamiliar interpersonal landscape transformed by a robot's presence. Don summarized his thoughts on the matter as such:

> *"You tend to focus on the individual rulings on the robot, but it's bigger than that. It's a family, it's a community of friends that potentially a robot could have contact with, and so that community has to sort of learn along with the robot and myself on what's acceptable and what's not acceptable. So it's a little bit bigger than just, like, how you define the rules, it's a little bit bigger and might change as time progresses and people understand how this works."*

Some participants took this concept of relational influence further, exploring how the responsibility for interacting with and maintaining a social robot's privacy system need not rest solely on its primary user. One impetus for this was the idea that it may not always be feasible for some older adults to operate a social robot completely on their own, no matter how generally usable its privacy system became. For Karen, this belief was informed by her own experience watching her mother age: *"When my mom got sick, the nature of the illness was that it affected her brain [...] she wouldn't have been able to do any of this."* Historic also stressed his witnessed experience with the negative effects of cognitive decline, disclosing that *"sometimes, knowing some circumstances of other people, is that they sort of hallucinate situations"* and act against their best interests, such as by *"[cutting] people off who should actually be part of the circumstances."*

To address these concerns, Karen, Bill, and Historic suggested that, in some cases, some people other than the main user should be allowed to exert a level of control over the robot's privacy system or even the robot as a whole. Karen and Bill determined that an individual with that level of power over the robot should be someone already trusted by the user, such as their power of attorney:

> Bill: *"You'd almost have to have someone with power of attorney—"*
> Karen: *"Controlling the robot or—"*
> Bill: *"Doing the things that normally the [user] would do."*

Historic similarly echoed this idea that such a large responsibility should fall on someone who has already had responsibility delegated to them through legal means: *"I think you just have to rely on the people that you've selected for your living will, people that are going to be executors of the estate, that they will have the best interests in mind."* Out of all the participants, Historic extended this concept of delegating control the furthest, with his suggestion in the access control design activity to grant *«authorized individuals to gain access»* to the robot with *«full control with directions in times of [user] incapacitation»*. Upon seeing the storyboard we developed based on this idea in the following session (see Figure 8), he elaborated that the process of escalating control and access privileges to another user should be tempered by a series of guardrails, such as *"a second party who would be a witness to the changes, especially on certain key things [. . .] that helps to kind of keep everybody in communication [. . .] [since] sometimes it might be that those two individuals have a difference in perspective."* In the end, Historic's elaboration serves as a reminder of the numerous and sometimes competing interests that often surround the management of an older adult's personal information, and the danger of taking an individualistic approach over one that has accounts for these relational dynamics.

5 Discussion

5.1 Design Implications

Our participants demonstrated a deep and meaningful interest in interacting directly with an access control system for a social robot. Their contributions indicate their willingness to engage in a number of administrative, intentional tasks, such as reviewing and auditing their robot's privacy policy, on a regular basis, as well as their desire to exercise 'total' control over their robot. For future designers, we offer a vision of ***the older adult user as their own privacy system administrator.*** Rather than expecting older adult users to be passive or apathetic, we suggest instead that social robots should be designed with an image of older adults as active and keen users who are eager to exert agency over their personal information.

One major design opportunity which follows from an understanding of older adults as active users of a robot's access control system is the enablement of intuitive privacy policy authoring for these users. Our participants' design outcomes provide evidence for the viability of a ***requirement delivery scheme based on CI factors.*** As discussed in Section 4.2, participants readily incorporated the CI factors of recipient and type of information when formulating their own rule structures, and generated multiple instantiations of relevant transmission principles on their own even when we did not explicitly instruct them to do so. Our work demonstrates a potential opportunity in exploring natural language policy authoring through CI factors for conversational agents in interpersonal settings and beyond.

Furthermore, their design outputs suggest a potential model for how policy adjustments could be made over time. The privacy requirements participants formulated tended to be both generally-worded and quite restrictive, at least initially. To use CPM parlance, we may think of a heavily restrictive access control policy as equivalent to a 'thick' privacy boundary shared between the user and their robot. We found that thick privacy boundaries were motivated by an initial feeling of wariness towards social robots and a sense of needing to be 'better safe than sorry' in the management of one's private information. However, privacy boundaries could become more 'thin' over time as a user developed trust in their robot or saw the utility in allowing their robot to reveal certain kinds of information to others. General or 'cover-all' privacy rules were mostly motivated insofar as they easily enabled blanket restrictions. Accounting for all of this, we tentatively propose that ***privacy requirements could be delivered to a social robot through a whitelisting or default-deny approach,*** where more specific privacy rules permitting the robot to reveal certain kinds of information to certain individuals under certain conditions could be added over time, with all other cases defaulting to a protective, easily understood baseline. While what we propose aligns conceptually with Saltzer and Schroeder's classic principle of least privilege [63], we wish to emphasize its socially normative dimension; as a social entity, a robot must earn the gradual trust of its user before it can come into full co-ownership of its user's privacy boundaries.

Overall, our participants preferred to author policies at a high level, conveying their privacy requirements through individual rules delivered in natural language. Translating natural language rules into machine-processable policies is an active area of research in the wider access control field [31]. However, participants also recognized that simply articulating privacy rules as a user in an unambiguous manner would be a challenge in and of itself. This is where the conversational nature of social robots could shine: as certain participants such as Don and Valerie imagined, a robot could ***alert its user to potential faults in the user's requirements and work to resolve them through dialogic processes.*** This process could, to an extent, mirror the dialogic interaction that occurs between two people when privacy boundary rules are explicitly being set.

A unique consideration for older adult users is that their life experience is often marked by gradual or sudden changes, whether physical, mental, or otherwise, which can catalyze vast shifts in interpersonal privacy dynamics. Thus, ***any access control system in our operational context must be ready to account for these changes.*** Dynamic and ever-changing access control requirements have design implications for the technical implementation of an access control system. For example, if policy verification is part of the system, verification processes may need to happen more often. Additionally, a policy authoring style that relies on the delivery of individual privacy rules which themselves may be subject to

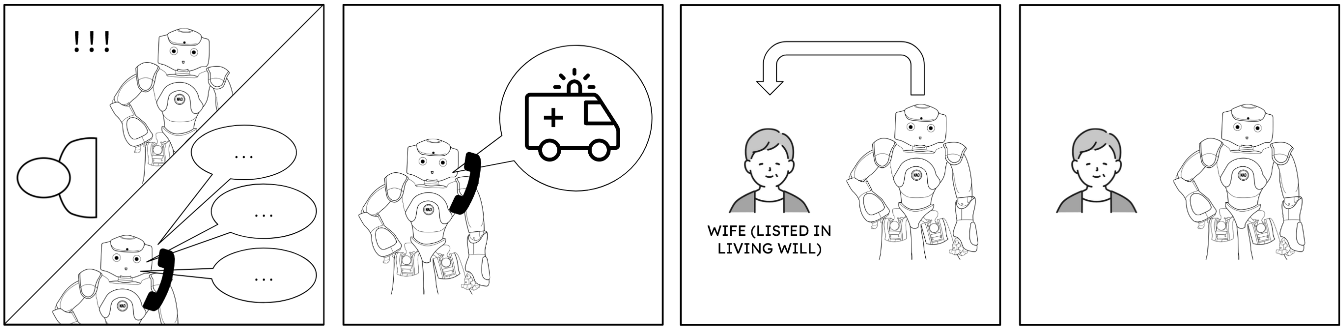

1) In the case of an event where I am incapacitated, my robot will act on its own to contact a list of my emergency contacts.

2) If the event is very severe, my robot will also call 911.

3) If I am incapacitated long enough, my robot will transfer full control over it to the individual(s) I have listed in my "living will."

4) This individual(s) has full access to every part of my robot. They can change privacy rules and access all of my information.

Figure 8: Storyboard prepared for Historic demonstrating his interest in the role of a robot and matters of control transferance during times of user incapacitation.

frequent revision will likely require a modular approach to low-level policy composition to ensure flexibility [9].

5.2 Relational Autonomy, Robots, and Privacy

Relational autonomy appeared as a somewhat unexpected yet highly salient theme in participant discussions. The term *relational autonomy* can be understood to refer to any number of approaches to the concept of autonomy that treat it as socially scaffolded rather than individualistic. Mackenzie and Stoljar defined relational autonomy as follows [45]:

> The term "relational autonomy," as we understand it [...] is rather an umbrella term, designating a range of related perspectives. These perspectives are premised on a shared conviction, the conviction that persons are socially embedded and that agents' identities are formed within the context of social relationships and shaped by a complex of intersecting social determinants, such as race, class, gender, and ethnicity. Thus the focus of relational approaches is to analyze the implications of the intersubjective and social dimensions of selfhood and identify for conceptions of individual autonomy and moral and political agency.

A relational treatment of autonomy seemed to come naturally to certain participants as they spoke to how an older adult's decisions for their social robot's privacy behaviour would be influenced by the older adult's context of social relationships. Participants discussed how an older adult's perception of their interests with regards to privacy boundary management would necessarily be shaped by the values and actions of the people closest to them: ideally, towards a mutually desirable end, as illustrated by Karen's explication of her narrative prototype.

Participants recognized the vulnerabilities faced by many older adults and readily speculated on the intersection of these vulnerabilities with relational, multi-user control over a single robot. Certain participants, such as Karen, Bill, and Historic, envisioned normative standards for the formal delegation of control over a robot from an older adult to another person. Participants brought up this operational possibility when discussing older adults with cognitive impairments or who would otherwise lack the capacities needed to operate their robot by themselves. When conceptualizing who might be involved in relational control of a social robot and its privacy system, participants looked to legal constructs such as powers of attorney and personal directives. The implications of this are multi-faceted. First, it should reinforce the idea that, should social robots become widely accepted sometime in the future, we ought to expect that they will play a highly integral and personal role in the lives of their users, enough to be indelibly intertwined with their users' expression of privacy and agency. Second, it suggests the potential future role legal policy could play in modelling new norms regarding social robots as respondents to a web of intersubjective demands.

We are interested in how a relational treatment of autonomy can come to inform more nuanced understandings of social robots as embedded within a social environment that gives rise to a user's judgement of and ability to exercise their autonomy. Much has been written about how social robots could instrumentally enhance or undermine their users' autonomy [21], but we believe there may be more insight to be gained if we explicitly include robots within the relational turn. A relational approach to ethics opens up new possibilities in reconsidering the moral standing of robots and our relationships with them [14]. Does a robot's physical embodiment and sociality owe it a form of 'real' social recognition [33]? If so, what does this have to say about its user's conception of their self and their autonomy, especially when the user is an older adult? By asking these questions, we suggest how future work can adopt relational ethics to explore what roles older adults desire for social robots, and how these roles will necessarily intertwine with older adults' autonomy and vulnerability [6].

5.3 Limitations and Future Work

Although a minority of our participants were affected by physical impairments, we do not deem any of them to have been affected by cognitive impairment during the time in which our co-design sessions were taking place. Our participants had active lifestyles and viewed themselves as generally healthy; a few remarked at various points during the sessions that they didn't see themselves as needing social robots or benefitting greatly from them at that moment in time. Our work addresses social robots generally, and takes pains to avoid positioning them exclusively within the category of 'elder care.' However, as certain themes and subsequent discussions in this paper have shown, it is impossible to fully disengage the 'care' aspect from gerontechnological social robot development. Furthermore, notions of care are inextricably linked to those of vulnerability, dependence, and relationality [44]. Therefore, a recommended conceptual extension of this work would be to involve a more focussed sample of older adults with greater care needs in a similar participatory research process. These participants may be able to shed further light on issues related to interpersonal privacy, vulnerability, and autonomy, as well as provide more concrete design suggestions or scenarios for social robots implicated in a salient context of care, such as when there is home care or the older adult user is living in a care facility. It may also be worth considering involving certain other stakeholders, such as formal or family caregivers, community workers, and healthcare providers in the co-design research process, a common tack in the gerontechnological co-design literature [47]. However, researchers who take this approach should continually and carefully examine whether they are privileging the accounts of secondary stakeholders over those of older adults, and be especially cognizant of the role of social and interpersonal power dynamics in shaping the outcomes of their work.

6 Conclusion

Social robots show promise for supporting the lives of older adults, but their privacy issues, especially in interpersonal social interactions, remain underexplored. Alongside six older adult participants, we conducted an exploratory co-design study pertaining to the design of access control-oriented privacy systems for social robots. We identified three overarching themes pertaining to various challenges and opportunities for the design of privacy-conscious social robots based on our participants' designerly speculations: *(1) Active Engagement with Privacy Systems; (2) Challenges and Opportunities for Translating Privacy Desires;* and *(3) Supporting Shifting and Social Dynamics*. We discussed high-level design implications following from these themes and our participants' specific contributions, and further discussed critical considerations surrounding older adults' autonomy and relationality against the backdrop of social robot use within privacy-sensitive contexts. Our work contributes early guidance for the design of privacy-conscious social robots for older adults, paving the way for future participant-informed design work in this gerontechnological space.

Acknowledgments

We graciously thank all of our participants: Bill, Don, Historic, John, Karen, and Valerie. We also thank Barry, who provided invaluable assistance with recruitment and pilot testing. Finally, we thank Jim Young for his helpful feedback on an earlier draft of this paper.

References

[1] Kenn Adams. 2007. *How to Improvise a Full-Length Play: The Art of Spontaneous Theater.* Allworth Press, New York.

[2] Ageless Innovation. 2018. JFA CAT PRODUCT MASTER AI. YouTube. (May 11, 2018). Retrieved July 22, 2025 from https://www.youtube.com/watch?v=mpULTJqPLbw.

[3] Ageless Innovation. 2024. Lifelike Robotic Pets for Seniors - Joy for All. Ageless Innovation LLC. Retrieved Feb. 24, 2024 from https://joyforall.com/.

[4] Zainab Agha, Zinan Zhang, Oluwatomisin Obajemu, Luke Shirley, and Pamela J. Wisniewski. 2022. A Case Study on User Experience Bootcamps with Teens to Co-Design Real-Time Online Safety Interventions. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. CHI '22: CHI Conference on Human Factors in Computing Systems. ACM, New Orleans, LA, USA, (Apr. 27, 2022), 40/1–8. doi:10.1145/3491101.3503563.

[5] Aldebaran. 2019. First Encounter with Pepper: an attentive medical companion. YouTube. (July 10, 2019). Retrieved July 22, 2025 from https://www.youtube.com/watch?v=XlDPMA6mn1g.

[6] Joel Anderson. 2013. Autonomy and Vulnerability Entwined. In *Vulnerability: New Essays in Ethics and Feminist Philosophy.* Studies in Feminist Philosophy. Catriona Mackenzie, Wendy Rogers, and Susan Dodds, (Eds.) Oxford University Press, New York, 134–161.

[7] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. 2018. Discovering Smart Home Internet of Things Privacy Norms Using Contextual Integrity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2, 2, (July 5, 2018), 59/1–23. doi:10.1145/3214262.

[8] Clara Berridge, Yuanjin Zhou, Amanda Lazar, Anupreet Porwal, Nora Mattek, Sarah Gothard, and Jeffrey Kaye. 2022. Control Matters in Elder Care Technology: Evidence and Direction for Designing It In. In *Designing Interactive Systems Conference.* DIS '22: Designing Interactive Systems Conference. ACM, Virtual Event Australia, (June 13, 2022), 1831–1848. doi:10.1145/3532106.3533471.

[9] Piero Bonatti, Sabrina De Capitani Di Vimercati, and Pierangela Samarati. 2000. A modular approach to composing access control policies. In *Proceedings of the 7th ACM Conference on Computer and Communications Security.* CCS00: 7th ACM Conference on Computer and Communication Security 2000. ACM, Athens, Greece, (Nov. 2000), 164–173. doi:10.1145/352600.352623.

[10] Virginia Braun and Victoria Clarke. 2013. *Successful Qualitative Research: A Practical Guide for Beginners.* SAGE, Los Angeles.

[11] Virginia Braun and Victoria Clarke. 2022. *Thematic Analysis: A Practical Guide.* SAGE, London; Thousand Oaks, California.

[12] Joost Broekens, Marcel Heerink, and Henk Rosendal. 2009. Assistive social robots in elderly care: a review. *Gerontechnology*, 8, 2, 94–103. doi:10.4017/gt.2009.08.02.002.00.

[13] L. F. Carver. 2020. Robopets: Using technology to monitor older adults raises privacy concerns. *The Conversation*, (Mar. 19, 2020). Retrieved May 26, 2026 from https://theconversation.com/robopets-using-technology-to-monitor-older-adults-raises-privacy-concerns-132326.

[14] Mark Coeckelbergh. 2018. What do we mean by a relational ethics?: Growing a relational approach to the moral standing of plants, robots, and other non-humans. In *Plant Ethics.* (1st ed.). Angela Kallhoff, Marcello Di Paola, and Maria Schörgenhumer, (Eds.) Routledge, London, 98–109.

[15] Sawyer Collins, Čedomir Stanojević, Casey Bennett, Zachary Henkel, Kenna Baugus Henkel, Nikki M. Abbott, Cindy L. Bethel, and Selma Šabanović. 2025. "Socially Assistive Robot Privacy Model": A Multi-model Approach to Evaluating Socially Assistive Robot Privacy Concerns. In *Social Robotics: 16th International Conference, ICSR + AI 2024.* Springer-Verlag, Odense, Denmark and Berlin, Heidelberg, 280–289. doi:10.1007/978-981-96-3522-1_25.

[16] Stephen Eick and Annie I. Anton. 2020. Enhancing Privacy in Robotics via Judicious Sensor Selection. In *2020 IEEE International Conference on Robotics and Automation (ICRA).* IEEE, Paris, France, (May 2020), 7156–7165. doi:10.1109/icra40945.2020.9196983.

[17] Katie Engelhart. 2021. What Robots Can—and Can't—Do for the Old and Lonely. *The New Yorker.* https://www.newyorker.com/magazine/2021/05/31/what-robots-can-and-cant-do-for-the-old-and-lonely.

[18] D. Feil-Seifer and M.J. Mataric. 2005. Socially Assistive Robotics. In *9th International Conference on Rehabilitation Robotics, 2005. ICORR 2005.* 9th International Conference on Rehabilitation Robotics, 2005. ICORR 2005. IEEE, Chicago, IL, USA, 465–468. doi:10.1109/ICORR.2005.1501143.

[19] Francisco Erivaldo Fernandes, Guanci Yang, Ha Manh Do, and Weihua Sheng. 2016. Detection of privacy-sensitive situations for social robots in smart homes. In *2016 IEEE International Conference on Automation Science and Engineering (CASE).* IEEE, Fort Worth, TX, (Aug. 2016), 727–732. doi:10.1109/coase.2016.7743474.

[20] Rachel L. Finn, David Wright, and Michael Friedewald. 2013. Seven Types of Privacy. In *European Data Protection: Coming of Age.* Serge Gutwirth, Ronald Leenes, Paul De Hert, and Yves Poullet, (Eds.) Springer Netherlands, Dordrecht, 3–32. doi:10.1007/978-94-007-5170-5_1.

[21] Paul Formosa. 2021. Robot Autonomy vs. Human Autonomy: Social Robots, Artificial Intelligence (AI), and the Nature of Autonomy. *Minds and Machines*, 31, 4, (Dec. 2021), 595–616. doi:10.1007/s11023-021-09579-2.

[22] FraunhoferIPA. 2012. Robotic home assistant Care-O-bot 3: possible application areas in everyday support. YouTube. (Dec. 17, 2012). Retrieved July 22, 2025 from https://www.youtube.com/watch?v=ABpOtvLzh2U.

[23] Susanne Frennert and Britt Östlund. 2014. Review: Seven Matters of Concern of Social Robots and Older People. *International Journal of Social Robotics*, 6, 2, (Apr. 2014), 299–310. doi:10.1007/s12369-013-0225-8.

[24] Global News. 2018. Meet Pepper – Canada's first talking, dancing, emotionally sensitive robot. YouTube. (May 7, 2018). Retrieved July 22, 2025 from https://www.youtube.com/watch?v=F6aDHtRFZJ0.

[25] David Gouaillier, Vincent Hugel, Pierre Blazevic, Chris Kilner, Jerome Monceaux, Pascal Lafourcade, Brice Marnier, Julien Serre, and Bruno Maisonnier. 2009. Mechatronic design of NAO humanoid. In *2009 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, Kobe, Japan, (May 2009), 769–774. doi:10.1109/ROBOT.2009.5152516.

[26] Reinhard Grabler and Sabine Theresia Koeszegi. 2024. Privacy beyond Data: Assessment and Mitigation of Privacy Risks in Robotic Technology for Elderly Care. *ACM Transactions on Human-Robot Interaction*, 14, 1, (Nov. 28, 2024), 6/1–23. doi:10.1145/3689216.

[27] Birgit Graf, Ulrich Reiser, Martin Hagele, Kathrin Mauz, and Peter Klein. 2009. Robotic home assistant Care-O-bot® 3 - product vision and innovation platform. In *2009 IEEE Workshop on Advanced Robotics and Its Social Impacts (ARSO)*. IEEE, Tokyo, Japan, (Nov. 2009), 139–144. doi:10.1109/ARSO.2009.5587059.

[28] Long-Jing Hsu, Philip B. Stafford, Weslie Khoo, Manasi Swaminathan, Kyrie Jig Amon, Hiroki Sato, Katherine M. Tsui, David J. Crandall, and Selma Sabanović. 2024. "Give it Time:" Longitudinal Panels Scaffold Older Adults' Learning and Robot Co-Design. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. HRI '24: ACM/IEEE International Conference on Human-Robot Interaction. ACM, Boulder, CO, USA, (Mar. 11, 2024), 283–292. doi:10.1145/3610977.3634979.

[29] Intuition Robotics. 2024. ElliQ, the sidekick for healthier, happier aging. ElliQ. Retrieved Feb. 24, 2024 from https://elliq.com/.

[30] Intuition Robotics. 2024. ElliQ. AI with EQ. YouTube. (Mar. 4, 2024). Retrieved July 22, 2025 from https://www.youtube.com/watch?v=1SkEu7LgTrE.

[31] Sakuna Harinda Jayasundara, Nalin Asanka Gamagedara Arachchilage, and Giovanni Russello. 2025. SoK: Access Control Policy Generation from High-level Natural Language Requirements. *ACM Computing Surveys*, 57, 4, (Apr. 30, 2025), 1–37. doi:10.1145/3706057.

[32] Somesh Jha, Ninghui Li, Mahesh Tripunitara, Qihua Wang, and William Winsborough. 2008. Towards Formal Verification of Role-Based Access Control Policies. *IEEE Transactions on Dependable and Secure Computing*, 5, 4, (Oct. 2008), 242–255. doi:10.1109/tdsc.2007.70225.

[33] Raya A. Jones. 2017. What makes a robot 'social'? *Social Studies of Science*, 47, 4, (Aug. 2017), 556–579. doi:10.1177/0306312717704722.

[34] Hyunjin Kang and Jeeyun Oh. 2023. Communication privacy management for smart speaker use: Integrating the role of privacy self-efficacy and the multidimensional view. *New Media & Society*, 25, 5, (May 2023), 1153–1175. doi:10.1177/14614448211026611.

[35] Myeung Un Kim, Harim Lee, Hyun Jong Yang, and Michael S. Ryoo. 2019. Privacy-Preserving Robot Vision with Anonymized Faces by Extreme Low Resolution. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Macau, China, (Nov. 2019), 462–467. doi:10.1109/iros40897.2019.8967681.

[36] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction*, 2, (Nov. 1, 2018), 102/1–31, CSCW, (Nov. 1, 2018). doi:10.1145/3274371.

[37] Hee Rin Lee, Selma Šabanović, Wan-Ling Chang, Shinichi Nagata, Jennifer Piatt, Casey Bennett, and David Hakken. 2017. Steps Toward Participatory Design of Social Robots: Mutual Learning with Older Adults with Depression. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. HRI '17: ACM/IEEE International Conference on Human-Robot Interaction. ACM, Vienna, Austria, (Mar. 6, 2017), 244–253. doi:10.1145/2909824.3020237.

[38] Tuck W. Leong and Benjamin Johnston. 2016. Co-design and Robots: A Case Study of a Robot Dog for Aging People. In *Social Robotics*. Vol. 9979. Arvin Agah, John-John Cabibihan, Ayanna M. Howard, Miguel A. Salichs, and Hongsheng He, (Eds.) Springer International Publishing, Cham, 702–711. doi:10.1007/978-3-319-47437-3_69.

[39] Lanjing Liu and Yaxing Yao. 2025. From Knowledge to Practice: Co-Designing Privacy Controls with Children. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. CHI 2025: CHI Conference on Human Factors in Computing Systems. ACM, Yokohama, Japan, (Apr. 26, 2025), 866/1–22. doi:10.1145/3706598.3713257.

[40] Yu-li Liu, Wenjia Yan, Bo Hu, Zhi Lin, and Yunya Song. 2024. Chatbots or Humans? Effects of Agent Identity and Information Sensitivity on Users' Privacy Management and Behavioral Intentions: A Comparative Experimental Study between China and the United States. *International Journal of Human–Computer Interaction*, 40, 19, (Oct. 2024), 5632–5647. doi:10.1080/10447318.2023.2238974.

[41] Y.-C. Liu, A. Chakrabarti, and T. Bligh. 2003. Towards an 'ideal' approach for concept generation. *Design Studies*, 24, 4, (July 2003), 341–355. doi:10.1016/s0142-694x(03)00003-6.

[42] Michal Luria, Rebecca Zheng, Bennett Huffman, Shuangni Huang, John Zimmerman, and Jodi Forlizzi. 2020. Social Boundaries for Personal Agents in the Interpersonal Space of the Home. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20: CHI Conference on Human Factors in Computing Systems. ACM, Honolulu HI USA, (Apr. 21, 2020), 1–12. doi:10.1145/3313831.3376311.

[43] Christoph Lutz, Maren Schöttler, and Christian Pieter Hoffmann. 2019. The privacy implications of social robots: Scoping review and expert interviews. *Mobile Media & Communication*, 7, 3, (Sept. 2019), 412–434. doi:10.1177/2050157919843961.

[44] Catriona Mackenzie. 2013. The importance of relational autonomy and capabilities for an ethics of vulnerability. In *Vulnerability: New Essays in Ethics and Feminist Philosophy*. Studies in Feminist Philosophy. Catriona Mackenzie, Wendy Rogers, and Susan Dodds, (Eds.) Oxford University Press, New York, 33–59.

[45] Catriona Mackenzie and Natalie Stoljar. 2000. Introduction: Autonomy Refigured. In *Relational Autonomy: Feminist Perspectives on Automony, Agency, and the Social Self*. Catriona Mackenzie and Natalie Stoljar, (Eds.) Oxford University Press, New York, 3–31.

[46] Nathan Malkin. 2023. Contextual Integrity, Explained: A More Usable Privacy Definition. *IEEE Security & Privacy*, 21, 1, (Jan. 2023), 58–65. doi:10.1109/MSEC.2022.3201585.

[47] Wendy Moyle, Deepa Sriram, Jenny Murfield, Lihui Pu, and Katarzyna Lion. 2025. Co-designed technologies and people living with dementia: A qualitative systematic review. *Contemporary Nurse*, (May 5, 2025). doi:10.1080/10376178.2025.2500377.

[48] Ivoline C. Ngong, Swanand Ravindra Kadhe, Hao Wang, Keerthiram Murugesan, Justin D. Weisz, Amit Dhurandhar, and Karthikeyan Natesan Ramamurthy. 2025. Protecting Users From Themselves: Safeguarding Contextual Privacy in Interactions with Conversational Agents. In *Findings of the Association for Computational Linguistics: ACL 2025*. Findings of the Association for Computational Linguistics: ACL 2025. Association for Computational Linguistics, Vienna, Austria, 26196–26220. doi:10.18653/v1/2025.findings-acl.1343.

[49] Helen Nissenbaum. 2019. Contextual Integrity Up and Down the Data Food Chain. *Theoretical Inquiries in Law*, 20, 1, (Mar. 16, 2019), 221–256. doi:10.1515/til-2019-0008.

[50] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review*, 79, 119–158.

[51] Anastasia K. Ostrowski, Cynthia Breazeal, and Hae Won Park. 2021. Long-Term Co-Design Guidelines: Empowering Older Adults as Co-Designers of Social Robots. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, Vancouver, BC, Canada, (Aug. 8, 2021), 1165–1172. doi:10.1109/RO-MAN50785.2021.9515559.

[52] Amit Kumar Pandey and Rodolphe Gelin. 2018. A Mass-Produced Sociable Humanoid Robot: Pepper: The First Machine of Its Kind. *IEEE Robotics & Automation Magazine*, 25, 3, (Sept. 2018), 40–48. doi:10.1109/MRA.2018.2833157.

[53] Jinkyung Katie Park, Renkai Ma, Naima Samreen Ali, Naulsberry Jean Baptiste, Zainab Agha, and Pamela J. Wisniewski. 2025. Teens, Privacy, and Algorithms: Navigating and Co-Designing Solutions for Interpersonal Boundary Management on Social Media. In *Proceedings of the 24th Interaction Design and Children*. IDC '25: Interaction Design and Children. ACM, Reykjavik, Iceland, (June 23, 2025), 589–607. doi:10.1145/3713043.3728840.

[54] Sandra Petronio. 2002. *Boundaries of Privacy: Dialectics of Disclosure. SUNY Series in Communication Studies*. State University of New York Press, (Oct. 10, 2002).

[55] Sandra Petronio. 2013. Brief Status Report on Communication Privacy Management Theory. *Journal of Family Communication*, 13, 1, (Jan. 2013), 6–14. doi:10.1080/15267431.2013.743426.

[56] Sandra Petronio, Jeffrey T. Child, and Robert D. Hall. 2021. Communication Privacy Management Theory: Significance for Interpersonal Communication. In *Engaging Theories in Interpersonal Communication*. (3rd ed.). Paul Schrodt, Kristina M. Scharp, and Dawn O. Braithwaite, (Eds.) Routledge, New York, 314–327.

[57] Anne Pfeifle. 2018. Alexa, What Should We Do about Privacy? Protecting Privacy for Users of Voice-Activated Devices. *Washington Law Review*, 93, 1, 421–458. https://digitalcommons.law.uw.edu/wlr/vol93/iss1/9/.

[58] Public Health Agency of Canada. 2025. Elder Abuse. Government of Canada. (Feb. 17, 2025). https://www.canada.ca/en/public-health/services/health-promotion/aging-seniors/elder-abuse.html.

[59] QualCoder. 2025. QualCoder. https://qualcoder.wordpress.com/.

[60] Wendy A. Rogers, Travis Kadylak, and Megan A. Bayles. 2022. Maximizing the Benefits of Participatory Design for Human–Robot Interaction Research With Older Adults. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 64, 3, (May 2022), 441–450. doi:10.1177/00187208211037465.

[61] Sea Rotmann. 2017. "Once upon a time…" Eliciting energy and behaviour change stories using a fairy tale story spine. *Energy Research & Social Science*, 31, (Sept. 2017), 303–310. doi:10.1016/j.erss.2017.06.033.

[62] Kavous Salehzadeh Niksirat, Evanne Anthoine-Milhomme, Samuel Randin, Kévin Huguenin, and Mauro Cherubini. 2021. "I thought you were okay": Participatory Design with Young Adults to Fight Multiparty Privacy Conflicts in Online Social Networks. In *Designing Interactive Systems Conference 2021*. DIS '21: Designing Interactive Systems Conference 2021. ACM, Virtual Event USA, (June 28, 2021), 104–124. doi:10.1145/3461778.3462040.

[63] J.H. Saltzer and M.D. Schroeder. 1975. The protection of information in computer systems. *Proceedings of the IEEE*, 63, 9, 1278–1308. doi:10.1109/PROC.1975.9939.

[64] Elizabeth B.-N. Sanders and Pieter Jan Stappers. 2008. Co-creation and the new landscapes of design. *CoDesign*, 4, 1, (Mar. 2008), 5–18. doi:10.1080/15710880701875068.

[65] Elizabeth B.-N. Sanders and Pieter Jan Stappers. 2012. *Convivial Design Toolbox: Generative Research for the Front End of Design*. BIS Publishers, Amsterdam.

[66] R.S. Sandhu and P. Samarati. 1994. Access control: principle and practice. *IEEE Communications Magazine*, 32, 9, (Sept. 1994), 40–48. doi:10.1109/35.312842.

[67] Shruti Sannon, Brett Stoll, Dominic DiFranzo, Malte F. Jung, and Natalya N. Bazarova. 2020. "I just shared your responses": Extending Communication Privacy Management Theory to Interactions with Conversational Agents. *Proceedings of the ACM on Human-Computer Interaction*, 4, (Jan. 4, 2020), 8/1–18, GROUP, (Jan. 4, 2020). doi:10.1145/3375188.

[68] Eli Saslow. 2026. To Stay in Her Home, She Let In an A.I. robot. *The New York Times*, (Feb. 12, 2026). Retrieved Apr. 12, 2026 from https://www.nytimes.com/2026/02/12/us/elliq-ai-robot-senior-companion.html.

[69] Andrea Slane and Isabel Pedersen. 2023. Older People's Perspectives on Privacy, Surveillance and Consumer Socially Assistive Robots: Paradoxes, Trade-offs and the Future of Privacy Protection. *SSRN Electronic Journal*. doi:10.2139/ssrn.4478713.

[70] Robert Sparrow and Linda Sparrow. 2006. In the hands of machines? The future of aged care. *Minds and Machines*, 16, 2, (Oct. 18, 2006), 141–161. doi:10.1007/s11023-006-9030-6.

[71] Terry Spencer. 2023. Chatty robot helps seniors fight loneliness through AI companionship. *AP News*, (Dec. 22, 2023). Retrieved Oct. 1, 2024 from https://apnews.com/article/artificial-intelligence-robot-elliq-senior-citizens-a343409477b7aea350254f94daf52eb7.

[72] Statistics Canada. 2019. Older adults and population aging statistics. Statistics Canada. (Feb. 8, 2019). Retrieved May 17, 2025 from https://www.statcan.gc.ca/en/subjects-start/older_adults_and_population_aging.

[73] Brian Tang, Dakota Sullivan, Bengisu Cagiltay, Varun Chandrasekaran, Kassem Fawaz, and Bilge Mutlu. 2022. Confidant: A Privacy Controller for Social Robots. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, Sapporo, Japan, (Mar. 7, 2022), 205–214. doi:10.1109/HRI53351.2022.9889540.

[74] Kristina A. Wenzel Egan and Colin Hesse. 2018. "Tell Me So That I Can Help You": Private Information and Privacy Coordination Issues in the Context of Eldercare. *Journal of Family Communication*, 18, 3, (July 3, 2018), 217–232. doi:10.1080/15267431.2018.1466784.

[75] Yaxing Yao, Justin Reed Basdeo, Smirity Kaushik, and Yang Wang. 2019. Defending My Castle: A Co-Design Study of Privacy Mechanisms for Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19: CHI Conference on Human Factors in Computing Systems. ACM, Glasgow Scotland Uk, (May 2, 2019), 198/1–12. doi:10.1145/3290605.3300428.

[76] Noé Zufferey, Kavous Salehzadeh Niksirat, Mathias Humbert, and Kévin Huguenin. 2024. Our Data, Our Solutions: A Participatory Approach for Enhancing Privacy in Wearable Activity Tracker Third-Party Apps. *Proceedings on Privacy Enhancing Technologies*, 2024, 4, (Oct. 2024), 734–754. doi:10.56553/popets-2024-0139.