Privacy at Home: an Inquiry into Sensors and Robots for the Stay at Home Elderly

Trenton Schulz[1], Jo Herstad[1], and Harald Holone[2]

[1] University of Oslo, Postbox 1072 Blindern, 0316 Oslo, Norway,
[trentonw|johe]@ifi.uio.no,
[2] Østfold University College, Postbox 700, 1757 Halden, Norway,
harald.holone@hiof.no

Abstract. The elderly in the future will use smart house technology, sensors, and robots to stay at home longer. Privacy at home for these elderly is important. In this exploratory paper, we examine different understandings of privacy and use Palen and Dourish's framework to look at the negotiation of privacy along boundaries between a human at home, the robot, and its sensors. We select three dilemmas: turning sensors on and off, the robot seeing through walls, and machine learning. We discuss these dilemmas and also discuss ways the robot can help make the elderly more aware of privacy issues and to build trust.

Keywords: robot, human-robot interaction, privacy, trust, elderly, home

1 Introduction

A popular solution to help older people stay at home longer is to use technology. Some examples of these solutions are *smart home* technology with different sensors around the house. These sensors detect and record different types of information: if someone is in the room, how much someone is breathing, measuring the pulse, etc. The information collected can be helpful as Goonawardene, Toh, and Tan [20] showed by using sensors to detect social isolation of seniors living at home.

These sensors in the house raise privacy issues for the elderly living at home. What are the sensors recording? Further, these sensors will typically be connected to the Internet and may send their data to other services to aid in machine learning to help future algorithms and robots. Who are they sending information to? Though many people may be unaware or indifferent to the sensors, placing all the sensors around the home may make the elderly feel they are under constant surveillance or they don't have any privacy at home. An older person may feel they have gained independence at the cost of being watched. Not to mention how visitors to the home may feel about the surveillance.

A different solution may be to introduce a robot into the home. The robot can carry many of these sensors and provide a mobile way of watching over the elderly person. A robot does not eliminate the need for sensors around the house since the robot may need stationary sensors to navigate and perform its duties. But it may offer a better way for the elderly person to relate to the sensors and understand the privacy issues involved with this technology. In addition, the robots could provide the elderly with an opportunity to negotiate their privacy since they could have an idea where the robot is instead of sensors that may be hidden around the house.

We are engaged in a research project investigating how a robot can assist elderly living at home. How is the potential privacy of a person affected by these sensors and robots in the home? What ways can the robot carry out its functions while preserving privacy? How can the robot help inform the elderly about privacy issues? How can we model these interactions between sensors, robots, and people? These are all general questions that need to be addressed. In this exploratory paper, we look at three dilemmas: turning sensors on and off, sensors that can see through walls, and machine learning. Our contribution is to highlight some privacy issues that appear when a robot is in the home of the elderly with different sensors, and make it possible to consider them in designing future human-robot interaction (HRI).

In the following paper, we will provide some background about the Internet of Things, privacy, and robots. We will use a privacy framework to show the boundaries of negotiation between the elderly at home, the robots and their sensors. We will present and discuss three dilemmas before discussing future directions and concluding the paper.

2 Background

In the following section, we will present background on the Internet of Things, smart homes, privacy, and robots. Though these concepts seem intuitive, it is useful to examine how they are used here.

2.1 The Internet of Things and the Smart Home

The Internet of Things (IoT) was originally meant as an idea for different things communicating in their own self-contained networks using technology like RFID [3]. But as devices and radios became more powerful and more energy efficient, it has changed to include the idea of items or “things” that communicate with each other over the Internet [4]. For example, wireless sensor networks that are used for monitoring patients [29] can be a form of the IoT. The *smart home* introduces concepts like ubiquitous computing using devices from the IoT.

As more devices are added that have a network connection outside the home, it is important to consider the privacy issues. Though each smart home is different, they likely contain devices and sensors that help the home do things or assist the people at home. For example, National Public Radio and Edison Research [30] looked at people owning speakers that are connected to online services like Amazon or Google (*smart speakers*) and found that people owning them tend to change habits to incorporate the speaker more into their daily routines. However, privacy issues have been highlighted by others [15] including the fact that they are always listening for commands.

A question that is raised when discussing sensors and robots is *trusting* them. Does the person trust the device to do what it should do and preserve the person's privacy? Schulz [41] looked at issues of creating trustworthy objects that can exist in smart homes, especially for people with disabilities or older people. Busch et al. [6] showed that the different requirements could be balanced to create a more trustworthy smart home artifact and that including people in the process helped create that result. The issue of trusting the device and assuming that people in the smart home can preserve their privacy has also been tested out in virtual and real smart home environments [7, 42]. Though focusing more on the IoT in general, Fritsch, Groven, and Schulz [19] discussed different strategies one could use when interacting with different items where one cannot determine if one can or should trust the item.

2.2 Privacy in a networked world

Privacy and technology have been an issue for decades. As Warren and Brandeis [46, p. 195] wrote:

> Instantaneous photographs and newspaper enterprise have invaded the sacred precincts of private and domestic life; and numerous mechanical devices threaten to make good the prediction that "what is whispered in the closet shall be proclaimed from the house-tops."

Yet it can be difficult to come up with an agreed upon definition of privacy. In performing a literature search in human-computer interaction (HCI), computer supported cooperative work (CSCW), and ubiquitous computing, Crabtree, Tolmie, and Knight [12] could not find any agreement on the concept of privacy. Rather they found understandings of privacy could be divided into several groups:

Control Privacy is understood as controlling the flow of personal information to who sees it. This is often attributed to Westin [47].

Boundary Privacy is understood as a set of boundaries that are negotiated by the person and who will have access. The basis for this understanding comes from Altman [1].

Contextual Integrity An understanding that is presented by Nissenbaum [31] as a flow of information that is appropriate based on the context. Others, like Ess and Fossheim [14] have used this understanding to show that privacy is evolving beyond to individual to include the idea of protecting relationships with others.

Paradox, Trade-off, and Concern Privacy is a paradox in that some people say one thing about privacy, but do another. Privacy can also drive concerns about what is being done with the personal information, and the trade-off of giving this information against its benefit.

Protective Measure Viewing entities that want access to personal information as an attack or threat, privacy is understood as protection against these attacks. This often focus on the technology and ways of implementing technologies [5]. Privacy by design [27] is also a part of this research. Another area of examination is the leaking of information through a *side channel*. That is, information leaked as a side effect of what you are doing [34]. Side channels can be technological or social.

All of these understandings of privacy are useful. In this paper, we will use privacy as negotiation based on boundaries (i.e., part of the Boundary understanding from above).

Building on the privacy work done by Altman [1], Palen and Dourish [33] proposed a framework that looks at privacy in a networked world. They identify three boundaries for negotiating privacy:

Disclosure boundary This boundary represents what you decide to tell others (disclose) and what you keep to yourself. Examples include: writing or speaking opinions about a subject in a public forum, placing posters on your lawn, and wearing clothing for a sports team.

Identity boundary This boundary represents the different roles we have in our lives. For example, in some areas we are an employee, other areas an enthusiast, and others a friend. Each of these roles have different kinds of commitments on what types of information we can and cannot share.

Temporal boundary This boundary represents persistence of information. This includes how information is handled over time and how building a history can reveal things about someone. The persistence of information often means you lose control over who and in what context information shows up at a later point in time. For example, knowing you made a phone call to a person versus knowing who, when, and the duration of all your phone calls.

There are other theories that have been built on top of Altman's work. For example, Petronio [35] developed Communication Privacy Management theory that creates rules for boundaries and disclosures. This is a more complex way of examining privacy issues, but Palen and Dourish [33] is more straight forward for this paper's explorations.

The framework has been used in contexts outside of the workplace examples in the original paper. For example, Holone and Herstad [21] used Palen and Dourish's framework to examine privacy issues in a social mapping application for accessibility issues. The framework highlighted the tension between keeping information private versus making it public. It also highlighted dangers of marking a workaround for accessibility as passable as an individual versus marking it passable as a representative of the handicap association since the latter may deter motivation to make a proper solution.

2.3 Robots and privacy

The term *robot* is used to describe things like algorithms, artificial intelligence, agents that live in a program (e.g., a *chatbot* on IRC or Slack), automated vehicles, or just "something new," especially when it has to deal with something that replaces a person. For this paper, we use the definition from the American

Heritage dictionary where a robot is defined as “a machine or device that operates automatically or by remote control,” [37]. That is, the paper looks at the machine, its sensors, and the software involved.

A robot has all the issues of other devices in the smart home and more. Robots need sensors to find their place in the environment or react to it. These sensors can gather different types of information, such as recording an image or audio. Kanda and Ishiguro [25] describes how many robots that interact with people (*social robots*) are dependent on sending the information from the sensors to other computers in a network to process the data and send responses back to the robot. The result is that the robot has more computing power than it would have due to its size or power constraints. But this transfer of information over the network can result in breaching the privacy of people in the area working with the robot.

Robots and privacy is a topic that is still being researched. Some early discussion of robots and privacy is from Kahn et al. [24] and Feil-Seifer, Skinner, and Matarić [17] who proposed privacy as one of the benchmarks for evaluating human-robot interaction. Syrdal et al. [44] interviewed people for a robot in a home scenario and asked what the robot should record. No one they interviewed was completely comfortable with a robot recording the information, but tolerated it if it was for an obvious purpose.

Young et al. [48] used social psychology to examine models for the acceptance of *domestic robots* (i.e., robots that are in the home, but are not necessarily communicated with socially). Young et al. noted that domestic robots would enter into personal spaces and deal with privacy issues. In the end, they found several factors that affect acceptance and perception of acceptance.

Calo [10] presented an overview of the privacy issues around surveillance and the fact that we act differently around anthropomorphic social robots. Later, Calo [11] posited that drones carrying cameras in public areas could make it easier for citizens to recognize the need for privacy.

A robot’s sensors and what they do may not be obvious. A study by Lee et al. [28] showed that people were not aware of the sensing capabilities of the robot (for example, that it could see behind itself) or a difference in what it collected and what it processed. This unawareness may even extend to standard cameras. Yet Caine, Šabanović, and Carter [9] ran an experimental study with a camera, a stationary robot, and a mobile robot to see how older adults changed their behavior to preserve their privacy. Caine, Šabanović, and Carter found that the older adults exhibited the most privacy-preserving behaviors when a camera was used and not a mobile robot. Schafer and Edwards [40] discussed this lack of transparency and other privacy (and copyright) issues. They argued that designs need to be more obvious for the people that will be using or interacting with the robot. Other projects [13, 2] have seen this need for respecting the privacy of an elderly person at home.

Some of the current research on robots and privacy has focused on robots that are operated by another person remotely and allows the person to be present and perform tasks in the environment where the robot is located (*telepresence* or *teleoperated robots*). The focus of this research is on obscuring the environment from the robot operator. Butler et al. [8] studied people's perceptions of privacy, and how well different video filters affected the operator's performance. Other types of filters have also proven effective [22, 23]. In a different type of experiment, Rueben et al. [38] used different interfaces for marking objects that should remain hidden to a robot's camera. In another experiment, Rueben et al. [39] found that informing a person who was operating the robot was important to the person's privacy concerns and what the robot did in the person's home.

Others have examined laws regarding privacy and robots. Pagallo [32] provides a summary of the different issues with privacy, Internet connectivity, robots, and what people can expect for privacy. Pagallo argues for the EU to examine these issues more. Fosch Villaronga and Roig [18] examined privacy issues with *carrier robots* (i.e., robots that carry people). They determined that the person has more control over the carrier robot than social robots, resulting in fewer privacy issues.

Eyssel, Wullenkord, and Nitsch [16] ran a study looking at interactions between a robot and whether the robot revealing something about itself and being likable would result in the human to reveal something as well. They found that it did not matter how much the participants liked the robot, but rather how much a participant anthropomorphized the robot.

In summary, there is a diverse and ongoing research in robots and privacy. We will take the framework from Palen and Dourish to examine elderly negotiate their privacy at home with a robot.

3 Examining Privacy Boundaries for the Elderly at Home

Palen and Dourish's framework can be used to examine the how the elderly negotiate privacy at home. Let's break this down into parts: the setting and devices—with a special look at the robot, the boundaries, and finally interactions between the boundaries and the connections.

3.1 The Home Setting

Imagine a person at home. This home has been outfitted with a robot and sensors for helping in monitoring the person. In addition, the person also has a wearable device and a smart phone that can monitor the person (Fig. 1). The person is elderly and there is risk of falling, the robot is not strong enough to help the person up, but it can signal others.

It will be rare for the person to only have a robot and no other connected devices. As mentioned in Sections 1 and 2.3, the robot is likely dependent on other sensors to help it find its way. It is also likely that either the robot or the sensors are connected to the Internet and exchanges data to aid in its monitoring work or as input to an algorithm to help the robot navigate around the home. The elderly person may also have extra devices around the house like smart speakers, phones, wearable devices, and other sensors. The interaction between the elderly person, the robot, its sensors, and other possible sensors and devices

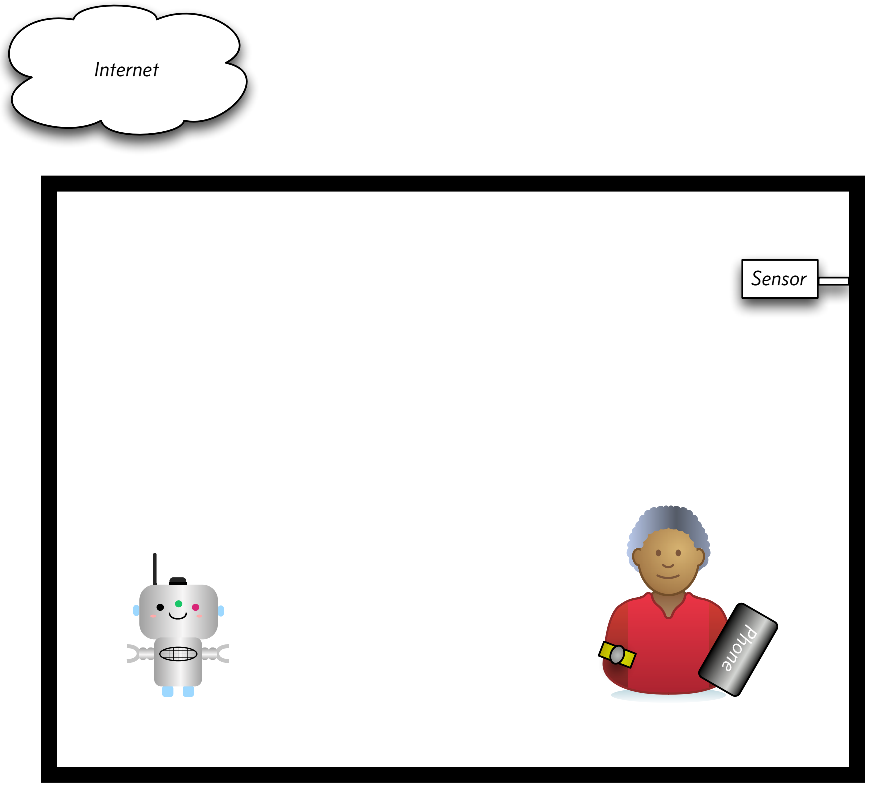

Fig. 1: Example set up for examining privacy issues in a smart home with a robot; the person also has a phone and a wearable device.

is interesting, it is also complex. For this paper, the focus will be on the elderly person, the robot, and its sensors and connections.

3.2 The Robot and Sensors

The example robot in this situation (Fig. 2) needs several kinds of sensors to perform its work, and it also has some ways of communicating. For simplicity, the sensors are all on the robot, but many of these sensors could also be used standalone in a smart home environment or the sensors could be mounted remotely and relay information back to the robot.

These sensors can be sorted by the two main senses: sight and sound. For sound, we have the standard microphone for listening to the person and the robot's environment. There may be more than one microphone to help the robot determine where the sound is coming from. The microphone may be an input device for human-robot interaction (i.e., the person can interact with the robot

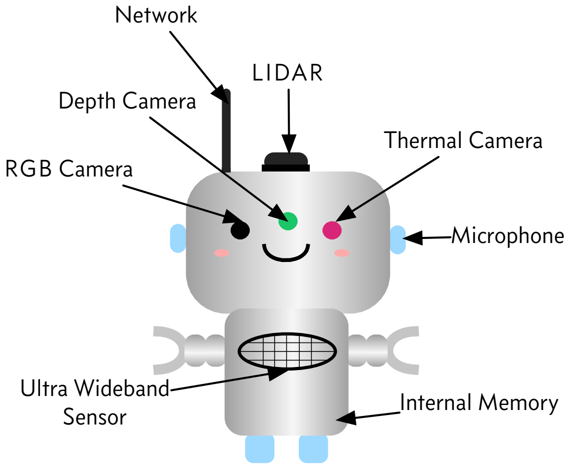

Fig. 2: Sensors that are on the robot in our model.

by speaking to it). The actual "listening" may not happen on the robot; it might be recorded by the microphone and then uploaded to another computer for processing like is done with smart speakers.

The next two technologies use radio waves. One is an ultra-wide band sensor as described by Tømmer, Kjelgård, and Lande [45]. These kind of sensors can be very sensitive and can be used to detect movement from several meters away (e.g., breathing or heart rate). These sensors can also detect this movement through walls. They can be thought of as special ears for detecting motion.

The second radio wave technology the robot has is a standard wireless connection (cellular, WI-FI, or both) that can be used for communicating with other networked devices (e.g., processing the sound above). This communication can get values from other sensors that may not be on the robot (e.g., the phone, wearable device, and sensor in the house in Fig. 4).

Our example robot is fitted with several kinds of eyes. One sensor is the Light Detection and Ranging device (or LIDAR for short). This sensor works by sending out a pulsed laser beam that can be used to measure the robot's distance from nearby objects. LIDAR is typically used to help a robot find its place in an area, and to notice things that may have moved since the last time it was there.

The example robot has three cameras. First, there is the RGB (Red, Green, Blue) camera, a regular camera that captures visible light. Second, there is a depth camera. Like LIDAR, this camera can show the distance between itself and an object. But LIDAR typically has a fixed height on a robot, while a depth camera will build depths of the entire scene it can see (for example, Fig. 3a). Finally, there is the thermal camera. This monitors infrared light to calculate the temperature of the objects it is pointed at.

Pictures from the cameras are useful and straightforward to interpret for people. But algorithms can analyze these images and determine where a person is in the system and create a skeleton for the person (Fig. 3b)

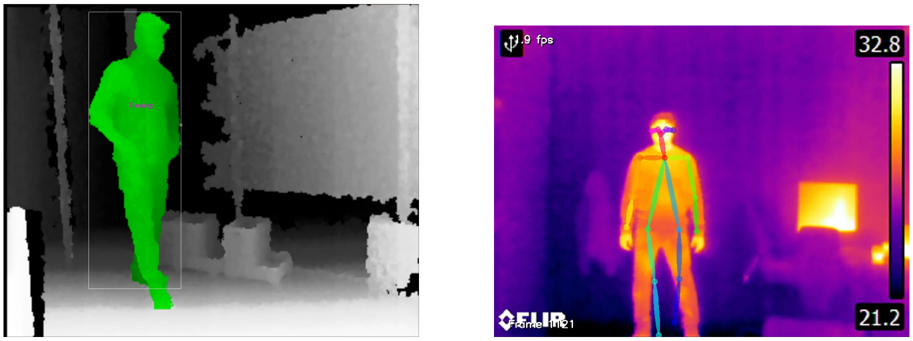

(a) Depth camera

(b) Thermal camera

Fig. 3: Example images from a depth camera (a) and a thermal camera (b).

Finally, there is the robot's storage and memory. The robot can record data from its sensors and use this data for later use. There is a limit to how much memory and storage a robot has, but the robot may also be linked to the Internet and the data may be stored some place else.

3.3 Negotiating Privacy

Palen and Dourish [33] boundaries help show areas where privacy is negotiated. Fig. 4 depicts boundaries of the elderly person, the robot, and the ways that that different conditions for negotiation. To summarize the boundaries:

Disclosure Boundary Represented by the dashed line in Fig. 4. In the home scenario, anything the elderly person says or does can reveal information about what is going on, what the person is doing, or where the person is. This includes who is in the house and what is said. In Fig. 4, the person negotiates the disclosure boundary when the person is sensed by the robot.

Identity Boundary Represented by the dotted line and also represented by the *ID* hat in Fig. 4. This is the identity of the person and is represented by a hat to emphasize that a person's role changes depending on the situation (for example, an elderly person may be a grandparent, friend, or president of the gardening club). This also shows that someone may hide other identities (here, hats you don't see) that only show up for certain occasions (for example, putting on the gardening club hat for a meeting). In Fig. 4, this is blurred with the disclosure boundary, but the elderly person negotiates the identity boundary with robot by deciding if the robot can detect the hat.

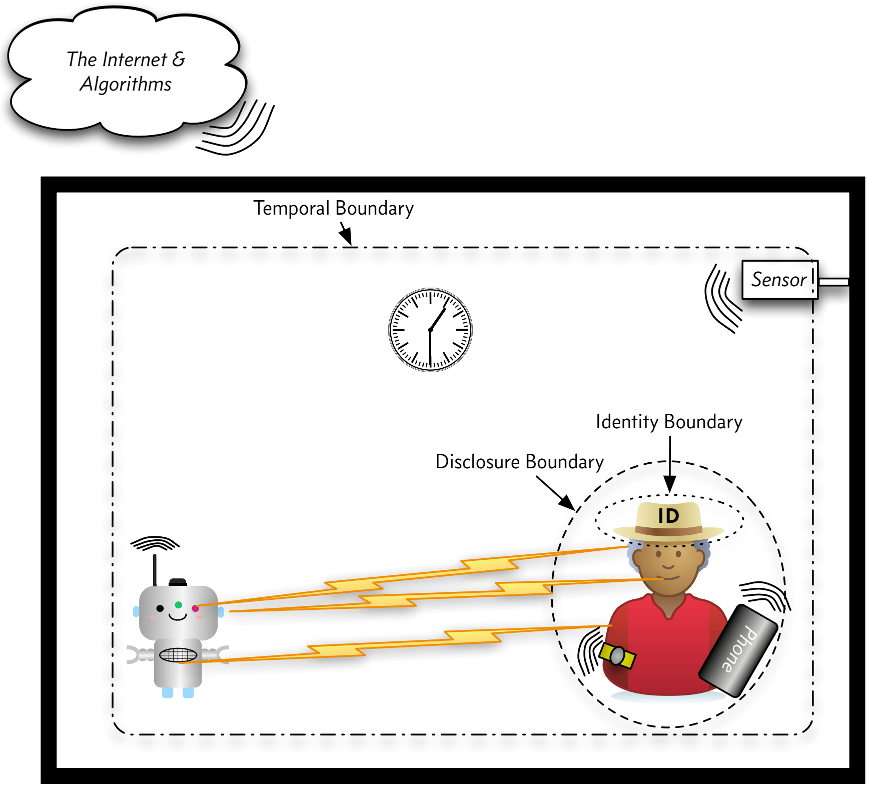

Fig. 4: The disclosure, identity, and temporal boundaries between the robot and the person along with the connections between the sensors.

Temporal Boundary Represented by dashed and dotted line going around the clock on the wall, the robot and the person in Fig. 4. This boundary is difficult to visualize in a static picture, but this represents time spent together with the robot and the person. A key factor here is also that the robot needs to remember the information from before. This also raises issues about why data is collected and what is to be used for. There is a need for the elderly person to negotiate with the temporal boundary when deciding what information the robot records and makes use of it later.

There are multiple ways to negotiate privacy with the robot. This is represented by the different (figurative) lightning bolts to and from the different parts of the human and robot. In addition, the different devices with the wave symbols can communicate with each other via a wireless radio; this includes the Internet represented by the cloud outside of the home.

Sometimes this negotiation is explicit between the robot, sensors, and the elderly. For example, the elderly person is in the same room as a robot. Other times, depending on how services work, the robot may end up doing these negotiations on the elderly person's behalf. For example, the robot understands speech by uploading recorded speech up to the Internet. Here, the robot has made the negotiation decision with the disclosure boundary to the Internet.

Negotiating on these boundaries (or the robot doing it on behalf of the person) is *not* intrinsically bad. The example robot is supposed to help with warning about the elderly person falling. To do that on its own, it needs to know where the person is and what the person is doing. This requires trust from the elderly that the robot's use of information will be responsible.

4 Dilemmas and Discussion

Let's examine some dilemmas that can come up with having a robot at home with the elderly. These three dilemmas help illustrate some privacy negotiations the person does with the robot. First, each dilemmas is presented. Next, the negotiation of privacy within this dilemma is presented. This is followed with some general discussion. For the first two dilemmas (Sections 4.1 and 4.2), only the disclosure and identity boundaries will be presented; the third dilemma's (Section 4.3) coverage of the temporal boundary is relevant for the other dilemmas.

4.1 Dilemma 1: Turning Sensors On and Off

Having control over the sensors is an obvious way of negotiating privacy. This means that the robot cannot detect the person while sensors are off. The robot may not be able to do much either. The connections would look like Fig. 1.

But it is not obvious how one can turn off all the controls in the system with robots and sensors. At the same time, the robot needs the sensors to perform its duties. How can the elderly person easily know which sensors are on and off to negotiate their privacy? The elderly person can ask the robot to turn on and off the sensors, but this requires trust that the robot will do the right thing.

Privacy Negotiation

Disclosure Boundary The negotiation here is the elderly person's desire for privacy. By turning on or off the sensors, the elderly person is explicitly stating a desire for privacy or willingness to share information with the robot.

Identity Boundary An elderly person's role may also be part of the negotiation with turning sensors on and off. For example, the person could be hosting visitors and doesn't want the robot in the way. This role of host is the reason the sensors are turned off. On the opposite side, a person could be ill or be temporarily disabled and need the robot on more than usual. This disabled role could also be deduced.

Discussion Turning on and off a robot's sensors is not different from other ways people negotiate surveillance in their everyday lives. For example, many people cover the cameras that are embedded in the displays of their laptops and phones. They remove these covers when they want to use the camera, but keep them covered otherwise. Even if the camera is on, it can only record darkness.

Another issue is the elderly person's desire for privacy versus the purpose the robot in the home. What happens when the elderly person needs the robot and cannot turn on the robot? Even if the robot can re-activate itself, if the person is out range of the robot, how can the robot find the person? How can it distinguish the person having left the house versus just being out of sensor range?

This balance between privacy and safety needs to be handled. This dilemma was also flagged by Amirabdollahian et al. [2] in their robot project. They suggested that the elderly person would have to figure out (i.e., negotiate) this balance. The elderly person would need to accept some loss of privacy in some situations. This has also been confirmed by discussions and interviews we've had with elderly. Some admit that they need help and are willing to have some sort of monitoring for this to happen. Another point is to figure out if turning off the sensors is recorded separately than other reasons for turning off the sensors (for example, a system reboot).

One way of finding this balance is to use a time limit for how long the sensors can be off—a *snooze*. In the situation where the robot is snoozing and the elderly need help, it may be a matter of waiting for the robot to wake up. This still leaves open the need to negotiate with all the other sensors if true privacy was desired. It may also be difficult to actually find and turn off all these sensors as well. Returning to the laptop example above, it is easy to cover the web camera, but it can be difficult to cover the accompanying microphones (if you can find them). This leads to the second dilemma.

4.2 Dilemma 2: Seeing Through Walls

Schafer and Edwards [40] put forth the idea that as long as you cannot see the robot, then the robot cannot see you. Schulz and Herstad [43] framed this as the idea of walking away from the robot. Of course, some elderly people may have issues with mobility, so the robot could also have the possibility of walking away, especially if it is asked.

This follows how we negotiate privacy with other people. If you can put something or some distance between you and the eyes of others (or the robot's sensors), you probably have some privacy. Fig. 4 shows that this walking away model could work. If the elderly's boundaries are out of range of the robot's cameras and microphones (or hidden from them), there is no interaction with the robot. Putting up a privacy wall between the robot and the elderly can accomplish this for cameras, and speaking softly may also circumvent the microphones. But if the robot has an ultra-wideband sensor that can sense people through the wall (Section 3.2), then the idea of negotiating privacy becomes difficult, and the walking away idea breaks down.

Privacy Negotiation

Disclosure Boundary The ultra-wideband sensors *eliminates* the chance for the person to negotiate presence (or lack of presence) in a room. This makes it difficult to trust the robot.

Identity Boundary The sensor may be able to pick up multiple people in an area and roles such as host could be picked up. If the elderly person is often in a room, it may also indicate a specific role the person has in that room.

Discussion The dilemma here is to balance the power of the sensors that can see through walls versus the elderly person's power to negotiate privacy.

There could be some negotiation on when the robot uses these sensors. For example, if the ultra-wide band sensor lets the robot sense through walls, perhaps the robot only uses this sensor only when it has difficulty finding the person? The robot could also do the equivalent of leaving the person alone once found. For example, by leaving or going to another room.

The robot could be designed so it is obvious when this sensor is in use. Another suggestion from Schafer and Edwards [40] is that robots sensors and what they are doing should be obvious to the elderly what they do. This could be a notification when it is using non-obvious sensors. This might be through sound, indicator lights, or some sort of other notification (for example like a recording light on a camera or an "on"-indicator on a microphone). This would help give the elderly person an opportunity to negotiate privacy instead of the robot overriding the process.

Instead of using the ultra-wide band sensors, one could choose sensors that can cut down in the amount of information leaked by the robot. For example, LIDAR is usually at a fixed height. As long as it can see its obstacles, it doesn't need to be much higher than the base of the robot. Pyo et al. [36] presented a prototype where a room contained a LIDAR and mirrors mounted just above floor level. They could monitor the room for robot navigation but not gather information about who people the person was in the room.

Another option is to use cameras that mask identifiable information in the images. Depth cameras and thermal cameras can show information about how someone is moving, but do not show details of the face like a regular camera does. They may also work better in situations where a regular camera does not (for example, low light). Algorithms can be trained to work with images from depth and thermal cameras as well. For example, Kido et al. [26] used thermal sensors to detect falls in toilets. Kido et al. motivation was to help preserve the privacy of elderly people, but still provide help if the person had fallen. There is negotiation with the disclosure boundary, but the amount of extra information the robot picks up is reduced.

4.3 Dilemma 3: Machine Learning

Many algorithms that robots use are based on machine learning. Suppose the robot in Fig. 4 watches the elderly person, and it uses its sensors to move safely around the house with the person. It also uses this information to determine when the person is moving regularly or differently and may fall in the future. This information is stored and used so that *other* robots may be further trained on noticing potential falls and moving around in a house. This is probably accepted as a positive thing. But what about when this data is then used by another algorithm to determine what rates to charge someone for a service? Or perhaps as a way of training a robot to capture people for law enforcement? This data has been used for a different person than what was originally decided and even though it may have been anonymized, the past actions of the person have helped train something to do something the person may not have been agreement with?

A similar issue is if the elderly person wants the robot to forget about something. It is possible to delete the specific instance, but what about if that data was used to train the machine learning algorithm? How does pulling the specific data out change the training data or how the algorithm developed from the previous data?

Privacy Negotiation

Temporal Boundary There's an interesting negotiation between machine-learning algorithms and the temporal boundary. Assuming that the robot must be *trained* or *learn* how it will interact with the elderly through watching the elderly and similar situations from others, the question becomes what happens if the data is used for training in other algorithms than the elderly person agreed to? In some ways, machine learning shows the implications of a negotiation with the temporal boundary on a large scale, going beyond the single person.

Disclosure and Identity Boundary There is no direct negotiation with the disclosure and identity boundaries. But the negotiation on the temporal boundary to collect more and more data to build algorithms may result in inferences that indirectly touch on these boundaries.

Discussion There is the negotiation with the temporal boundary and the use of previous information to help the robot perform its work. Since machine learning is built up from data from previous encounters, it's difficult for it to work without having access to previous information.

There is an obvious technical solution to this problem. Just have the robot delete the data that it has from time-to-time. This removes the direct negotiation with the temporal boundary, but there is the indirect negotiation through machine learning. Ideally, the data should be removed from training data and the algorithm retaught, but this may be impossible. Likely more data will be added to keep improving the algorithm. So, in some ways, the original temporal data may be gone, but it won't be forgotten in the algorithms.

Ultimately, this may be something that must be controlled through legal and technical solutions. This may make the temporal barrier a more important barrier to negotiate for the elderly, but it doesn't solve the problem of removing the data forever. A strict negotiation on the temporal boundary may result that the robot may never "get to know" the elderly person in the home.

5 Conclusion and Future Work

This exploratory look at robots and sensors at home with the elderly has shown that there are several privacy issues that need to be considered. We've also seen that the elderly need to be aware of the recording so they have an opportunity to negotiate their privacy with the robots and sensors. We presented three dilemmas. Using Palen and Dourish's framework, we could show tensions between providing help for the elderly at home and negotiating their privacy. There is no one-size-fits-all answer for preserving privacy and having the elderly live safely at home. As the dilemmas show, it's an individual negotiation.

Beyond privacy, there's also trust issues with the robot. The robot may be designed so that it is pleasing to the eye and easier to accept in the home than a sensor in a wall. But the dilemmas above gives reason to pause in trusting the robot: is it a helpful sheep that watches you or a wolf in sheep's clothing?

A robot needs to be trusted by the elderly. This goes beyond pleasing appearance and movement. The robot should be designed with protecting the privacy of the elderly and the elderly must have an opportunity to negotiate their privacy.

Though this exploratory work has found potential dilemmas, it is also important to get experiences from actual elderly at home. We have begun conducting interviews with the elderly living at home and are analyzing the transcripts. We will use the insights we gain from the analysis to develop prototype robots to help the elderly negotiate privacy at home and provide service that they value.

Acknowledgments This work is partly supported by the Research Council of Norway as a part of the Multimodal Elderly Care Systems (MECS) project, under grant agreement 247697. Thanks also to Md Zia Uddin for help in providing pictures from depth and thermal cameras.

References

1. Altman, I.: The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding. (1975/00/00)
2. Amirabdollahian, F., op den Akker, R., Bedaf, S., Bormann, R., Draper, H., Evers, V., Gelderblom, G., Gutierrez Ruiz, C., Hewson, D., Ninghang Hu, Iacono, I., Koay, K., Krose, B., Marti, P., Michel, H., Prevot-Huille, H., Reiser, U., Saunders, J., Sorell, T., and Dautenhahn, K.: Accompany: Acceptable robotiCs COMPanions for AgeiNG Years - Multidimensional Aspects of Human-System Interactions. In: 2013 6th International Conference on Human System Interactions (HSI), pp. 570–577. IEEE (2013)
3. Ashton, K.: That 'Internet of Things' Thing. RFID Journal (2009)
4. Atzori, L., Iera, A., and Morabito, G.: The Internet of Things: A Survey. Computer Networks 54(15), 2787–2805 (2010)

5. Bellotti, V., and Sellen, A.: Design for Privacy in Ubiquitous Computing Environments. In: Proceedings of the Third European Conference on Computer-Supported Cooperative Work 13–17 September 1993, Milan, Italy ECSCW '93. Ed. by G. de Michelis, C. Simone, and K. Schmidt, pp. 77–92. Springer Netherlands(1993)
6. Busch, M., Hochleitner, C., Lorenz, M., Schulz, T., Tscheligi, M., and Wittstock, E.: All In: Targeting Trustworthiness for Special Needs User Groups in the Internet of Things. In: Trust and Trustworthy Computing. Ed. by M. Huth, N. Asokan, S. Čapkun, I. Flechais, and L. Coles-Kemp, pp. 223–231. Springer Berlin Heidelberg(2013)
7. Busch, M., Lorenz, M., Tscheligi, M., Hochleitner, C., and Schulz, T.: Being There For Real – Presence in Real and Virtual Environments and Its Relation to Usability. In: Proceedings of the 8th Nordic Conference on Human-Computer Interaction Fun, Fast, Foundational - NordiCHI '14, pp. 117–126. ACM Press, New York, New York, USA (2014)
8. Butler, D.J., Huang, J., Roesner, F., and Cakmak, M.: The Privacy-Utility Tradeoff for Remotely Teleoperated Robots. In: Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction. HRI '15, pp. 27–34. ACM, New York, NY, USA (2015)
9. Caine, K., Šabanović, S., and Carter, M.: The Effect of Monitoring by Cameras and Robots on the Privacy Enhancing Behaviors of Older Adults. In: 2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pp. 343–350 (2012)
10. Calo, R.: Robots and Privacy. SSRN Scholarly Paper ID 1599189, Rochester, NY: Social Science Research Network (2010)
11. Calo, R.: The Drone as Privacy Catalyst. SSRN Scholarly Paper ID 2340753, Rochester, NY: Social Science Research Network (2011)
12. Crabtree, A., Tolmie, P., and Knight, W.: Repacking 'Privacy' for a Networked World. Computer Supported Cooperative Work (CSCW) 26, 453–488 (2017)
13. Draper, H., and Sorell, T.: Ethical Values and Social Care Robots for Older People: An International Qualitative Study. Ethics Inf Technol 19(1), 49–68 (2017)
14. Ess, C., and Fossheim, H.: Personal Data: Changing Selves, Changing Privacies. In: Digital Enlightenment Yearbook 2013, pp. 40–55. IOS Press(2013)
15. Estes, A.C.: Don't Buy Anyone an Echo. Gizmodo (2017)
16. Eyssel, F., Wullenkord, R., and Nitsch, V.: The Role of Self-Disclosure in Human-Robot Interaction. In: 2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 922–927 (2017)
17. Feil-Seifer, D., Skinner, K., and Matarić, M.J.: Benchmarks for Evaluating Socially Assistive Robotics. Interaction Studies 8(3), 423–439 (2007)
18. Fosch Villaronga, E., and Roig, A.: European Regulatory Framework for Person Carrier Robots. Computer Law & Security Review 33(4), 502–520 (2017)
19. Fritsch, L., Groven, A.-K., and Schulz, T.: On the Internet of Things, Trust Is Relative. In: Constructing Ambient Intelligence. Ed. by R. Wichert, K. Laerhoven, and J. Gelissen, pp. 267–273. Springer Berlin Heidelberg, Berlin(2012)
20. Goonawardene, N., Toh, X., and Tan, H.-P.: Sensor-Driven Detection of Social Isolation in Community-Dwelling Elderly. In: Human Aspects of IT for the Aged Population. Applications, Services and Contexts. LNCS, pp. 378–392. Springer, Heidelberg (2017)
21. Holone, H., and Herstad, J.: Negotiating Privacy Boundaries in Social Applications for Accessibility Mapping. In: Proceedings of the 6th Nordic Conference on Human-

Computer Interaction: Extending Boundaries. NordiCHI '10, pp. 217–225. ACM, New York, NY, USA (2010)

22. Hubers, A., Andrulis, E., Scott, L., Stirrat, T., Zhang, R., Sowell, R., Rueben, M., Grimm, C.M., and Smart, W.D.: Using Video Manipulation to Protect Privacy in Remote Presence Systems. In: Social Robotics. Ed. by A. Tapus, E. André, J.-C. Martin, F. Ferland, and M. Ammi, pp. 245–254. Springer International Publishing(2015)
23. Hubers, A., Andrulis, E., Smart, W.D., Scott, L., Stirrat, T., Tran, D., Zhang, R., Sowell, R., and Grimm, C.: Video Manipulation Techniques for the Protection of Privacy in Remote Presence Systems. In: Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts. HRI'15 Extended Abstracts, pp. 59–60. ACM, New York, NY, USA (2015)
24. Kahn Jr., P.H., Ishiguro, H., Friedman, B., Kanda, T., Freier, N.G., Severson, R.L., and Miller, J.: What Is a Human?: Toward Psychological Benchmarks in the Field of Human–robot Interaction. Interaction Studies 8(3), 363–390 (2007)
25. Kanda, T., and Ishiguro, H.: Human-Robot Interaction in Social Robotics. CRC Press (2012)
26. Kido, S., Miyasaka, T., Tanaka, T., Shimizu, T., and Saga, T.: Fall Detection in Toilet Rooms Using Thermal Imaging Sensors. In: 2009 IEEE/SICE International Symposium on System Integration (SII), pp. 83–88 (2009)
27. Langheinrich, M.: Privacy by Design — Principles of Privacy-Aware Ubiquitous Systems. In: Ubicomp 2001: Ubiquitous Computing. Ed. by G. Abowd, B. Brumitt, and S. Shafer, pp. 273–291. Springer Berlin / Heidelberg(2001)
28. Lee, M.K., Tang, K.P., Forlizzi, J., and Kiesler, S.: Understanding Users' Perception of Privacy in Human-Robot Interaction. In: Proceedings of the 6th International Conference on Human-Robot Interaction. HRI '11, pp. 181–182. ACM, New York, NY, USA (2011)
29. Leister, W., Schulz, T., Lie, A., Grythe, K., and Balasingham, I.: Quality of Service, Adaptation, and Security Provisioning in Wireless Patient Monitoring Systems. In: Biomedical Engineering, Trends in Electronics, Communications and Software. Ed. by A.N. Laskovski, pp. 711–736. InTech(2011)
30. National Public Radio, and Edison Research: The Smart Audio Report: Fall–Winter 2017, p. 31. National Public Radio (2017)
31. Nissenbaum, H.: Privacy as Contextual Integrity. Washington Law Review 79(1), 119–157 (2004)
32. Pagallo, U.: Robots in the Cloud with Privacy: A New Threat to Data Protection? Computer Law and Security Review 29(5), 501–508 (2013)
33. Palen, L., and Dourish, P.: Unpacking "Privacy" for a Networked World. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. CHI '03, pp. 129–136. ACM, New York, NY, USA (2003)
34. Percival, C.: Some Thoughts on Spectre and Meltdown, (2018). `http://www.daemonology.net/blog/2018-01-17-some-thoughts-on-spectre-and-meltdown.html` (visited on 01/31/2018)
35. Petronio, S.S.: Boundaries of Privacy: Dialectics of Disclosure. State University of New York Press, Albany (2002)
36. Pyo, Y., Hasegawa, T., Tanaka, M., Tsuji, T., Morooka, K., and Kurazume, R.: Measurement and Estimation of Indoor Human Behavior of Everyday Life Based on Floor Sensing with Minimal Invasion of Privacy. In: pp. 2170–2176. IEEE (2013)
37. *Robot.* In: American Heritage Dictionary of the English Language. Houghton Mifflin Harcourt Publishing Company(2011)

38. Rueben, M., Bernieri, F.J., Grimm, C.M., and Smart, W.D.: Evaluation of Physical Marker Interfaces for Protecting Visual Privacy from Mobile Robots. In: 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 787–794 (2016)
39. Rueben, M., Bernieri, F.J., Grimm, C.M., and Smart, W.D.: Framing Effects on Privacy Concerns About a Home Telepresence Robot. In: Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction. HRI '17, pp. 435–444. ACM, New York, NY, USA (2017)
40. Schafer, B., and Edwards, L.: "I Spy, with My Little Sensor": Fair Data Handling Practices for Robots between Privacy, Copyright and Security. Connection Science 29(3), 200–209 (2017)
41. Schulz, T.: Creating Universal Designed and Trustworthy Objects for the Internet of Things. In: Learning and Collaboration Technologies. Technology-Rich Environments for Learning and Collaboration. Ed. by P. Zaphiris and A. Ioannou, pp. 206–214. Springer International Publishing(2014)
42. Schulz, T., Fuglerud, K.S., Arfwedson, H., and Busch, M.: A Case Study for Universal Design in the Internet of Things. In: Universal Design 2014: Three Days of Creativity and Diversity. Ed. by H. Caltenco, P.-O. Hedvall, A. Larsson, K. Rassmus-Gröhn, and B. Rydeman, pp. 45–54. IOS Press(2014)
43. Schulz, T., and Herstad, J.: Walking Away from the Robot: Negotiating Privacy with a Robot. In: Proceedings of the 31th International BCS Human Computer Interaction Conference (HCI 2017). BCS Learning and Development, Sunderland, UK (2018)
44. Syrdal, D.S., Walters, M.L., Otero, N., Koay, K.L., and Dautenhahn, K.: He Knows When You Are Sleeping-Privacy and the Personal Robot Companion. In: Proc. Workshop Human Implications of Human–robot Interaction, Association for the Advancement of Artificial Intelligence (AAAI'07), pp. 28–33 (2007)
45. Tømmer, M., Kjelgård, K.G., and Lande, T.S.: Body Coupled Wideband Monopole Antenna. In: 2016 Loughborough Antennas Propagation Conference (LAPC), pp. 1–5 (2016)
46. Warren, S.D., and Brandeis, L.D.: The Right to Privacy. Harvard Law Review 4(5), 193–220 (1890)
47. Westin, A.: Privacy and Freedom. Atheneum, New York (1967)
48. Young, J.E., Hawkins, R., Sharlin, E., and Igarashi, T.: Toward Acceptable Domestic Robots: Applying Insights from Social Psychology. Int J of Soc Robotics 1(1), 95 (2009)