Privacy Agents: Utilizing Large Language Models to Safeguard Contextual Integrity in Elderly Care

Reinhard Grabler*
reinhard.grabler@tuwien.ac.at
TU Wien
Vienna, Austria

Matthias Hirschmanner†
matthias.hirschmanner@tuwien.ac.at
TU Wien
Vienna, Austria

Helena Anna Frijns*
helena.frijns@tuwien.ac.at
TU Wien
Vienna, Austria

Sabine Theresia Koeszegi*
sabine.koeszegi@tuwien.ac.at
TU Wien
Vienna, Austria

ABSTRACT

A value-based design process in the development of robotic technologies for elderly care requires approaches to protect privacy. With the rise of Large Language Models (LLMs) new use cases for robotic technology can be facilitated. In this paper we present a conceptual approach to utilizing LLMs to enhance privacy. We refer to a use case of a care documentation support agent that should aid care workers in their care routines. Our contribution is based on the understanding of Nissenbaum's privacy as contextual integrity. We introduce a privacy agent that continuously monitors information flows of recorded conversations, and identifies key parameters that are compared to the contextual privacy norms to detect privacy violations.

CCS CONCEPTS

• **Security and privacy** → **Privacy protections**; • **Computer systems organization** → *Robotics*.

KEYWORDS

Privacy, Contextual Integrity, LLM, Elderly Care, Robotics

1 INTRODUCTION

The development of robotic technology in care is an important research direction for the future. With a demographic shift ahead, there will be more older adults in need of care, while the number of potentially available caregivers for these older adults will fall; there will be fewer young people to fill this gap as informal and professional caregivers [31, 34]. Finding suitable roles for robotic technology in care is a task that requires trans-disciplinary research. In the area of tension between the professional identity of caregivers, institutional control, and the value of good care, it is important to pursue value-based design. Introducing robots into sensitive contexts, with sensors, and the ability to move and interact with people, is a particular challenge for the value of privacy (e.g., [11, 21]). Nevertheless, there are many approaches to mitigate privacy concerns in robotic technology for elderly care. Those approaches range from strictly regulating data collection [23] to engineering solutions such as shape extraction [12], blurring video [33], or robot behaviors such as turning away cameras [14]. Nevertheless, the multi-dimensional nature of privacy [22] calls for innovative solutions. The concept of privacy as contextual integrity, formulated by Nissenbaum [27], understands privacy as contextually bound. Privacy violations occur if parameters of an information flow do not align with the contextual norms.

With the recent rise of Large Language Models (LLMs), we see a potential to identify the parameters within human-to-human conversations and decide whether it is privacy appropriate for a robotic agent to record and process this information. In this publication, we showcase how LLMs can be facilitated to negotiate privacy in the context of care. We present a use case for a documentation support agent in the institutional care context and sketch a conceptual framework of a *privacy agent*, which helps to identify the parameters of an information flow and compares it to existing contextual privacy norms.

2 THEORETICAL BACKGROUND

Privacy is a term which is not conclusively defined. A commonly referred to definition of privacy is the fundamental "right of the individual to be let alone" [9]. Altman [3] on the other hand understands privacy as a dynamic process where boundaries are established and interactions regulated. The authors Burgoon [10] and Leino-Kilpi et al. [22] conceptualize privacy as a multi-dimensional framework along the informational, physical, social, and psychological dimensions. Another approach to privacy is the understanding of "Contextual Integrity," as it was formulated by Nissenbaum [27], where violations of privacy are context-dependent.

2.1 Privacy Impact of Robots in Care

Privacy concerns could have a negative effect on the adoption of robotic technologies in care [2, 6]. Nevertheless, the impact on privacy varies depending on the type of robot introduced into a context. When considering social robots, the four dimensional privacy view (informational, physical, social, psychological) is particularly useful. Social bonds might form between humans and robots, thus affecting

*Authors are with the Institute of Management Science.
†Author is with the Automation and Control Institute.

Privacy-Aware Robotics Workshop, HRI '24, March 11–15, 2024, Boulder, CO, USA

dimensions of psychological and social privacy. Furthermore, robots impact aspects of physical privacy due to their mobility, and they collect and process vast amounts of data, leading to informational privacy concerns [23]. According to Calo [11], introducing mobile robots is a novel challenge for privacy where neither law nor technology can prevent harms, affecting the categories of surveillance, increased access and social meaning. Yusif et al. [39] name privacy as a top concern of older adults for the willingness of adoption of assistive technologies. Especially situations where nakedness is involved are an issue [14, 38].

There is also empirical evidence that it makes a difference who is on the receiving end of information collected with robotic technology. A study [15] identified the requirements that only authorized persons should have access, that security is important, and that only useful information is stored. Moreover, in care settings healthcare personnel tends to be trusted more with health information than relatives or other third parties [13].

Nevertheless, not all studies confirm the importance of privacy in the adoption of robotic technology in care. Lutz and Tamó-Larrieux [26] find that there are only moderate concerns about the privacy impact posed by social robots, pointing out the privacy paradox, where the perceived benefits of the technology outweigh the privacy issues. Another study [17] points out that older adult participants in their experiment were not worried about disclosing information to a robotic agent unless they were repeatedly told that privacy should be a concern.

2.2 Privacy as Contextual Integrity

The core concept of the framework of contextual integrity is that privacy norms vary for different contexts [27]. For the context of care, it might be appropriate to reveal health information to healthcare personnel, but not to bystanders. Thus, perceived privacy depends on the social roles of the participants, as well as on the type of information that is disclosed. Privacy, in Nissenbaum's view, can then be understood as a flow of personal information, for which context-relative informational norms exist with the following key parameters [4, 28, 37]:

(1) *Sender*: Who is sending the information?
(2) *Recipient*: Who is on the receiving end of the information?
(3) *Attribute*: What kind of information is transmitted?
(4) *Subject*: Whom is the information about; whom does it affect?
(5) *Transmission principle*: What are the conditions/constraints for the transmission, e.g., is it confidential or not?

When one of the parameters in the information flow does not meet the applicable contextual norms, this constitutes a privacy violation [4]. With this approach, robotic technology should adhere to the contextual norms, instead of blindly collecting all data [24]. This could reduce the impact of robots on the multiple dimensions of privacy (see Section 2.1). If the participants in a context take note of the robot as an agent that is capable of differentiating whether it is appropriate to, e.g., video record, in a certain situation, concerns regarding physical privacy are reduced. The same can be applied to social and psychological privacy, e.g., if the robot does not record a secret that it is told by the user. While the concept of contextual integrity is commonly referred to as a promising solution for mitigating privacy concerns posed by robotic technology in care (e.g., [23–25, 32, 35, 36]), there is little evidence regarding its practical implementation, as it is a novel frontier for the domain of Human–Robot Interaction (HRI) [33].

3 CONCEPTUAL FRAMEWORK

The conceptual framework for privacy-aware robotics we present in this publication is based on the theoretical background of privacy as contextual integrity (see Section 2.2). In the course of a series of participatory workshops, a use case for LLMs emerged, which forms the basis of the present conceptual framework for preserving privacy by means of contextual integrity. In order to better picture the application, the experimental use case is outlaid.

3.1 Use Case: Care Documentation Support

In order to co-design robotic technology with stakeholders from the care sector, we conducted a participatory workshop series of 5 workshops with 13 care workers and 12 residents in 2 Austrian residential care homes. Several aspects of key robotic technologies were introduced, among them the use of LLMs as conversational agents. One theme that appeared repeatedly in the workshops from the side of the care workers was the need for supporting documentation. When introduced to the possibilities of LLMs, they saw potential. Consequently, we developed a prototype for a documentation support agent. Although this software agent can potentially operate on a robot, we currently utilize a microphone and process the audio on a workstation computer. In an experimental setups we recorded audio in a care scenario, which was staged by care workers with the "Nursing Anne" training mannequin. We transcribed the recorded audio with WhisperX [7] using Whisper's `large-v3` model by OpenAI [30]. As the spoken language is German, we used the EM German model [16], an openly available LLM based on Mistral7B [20]. The task of the documentation support agent is to summarize the care tasks done by the care workers, in order to assist them in their documentation later. The text summarized by the LLM is not directly entered into the documentation system, but is reviewed by care workers beforehand to ensure human oversight.

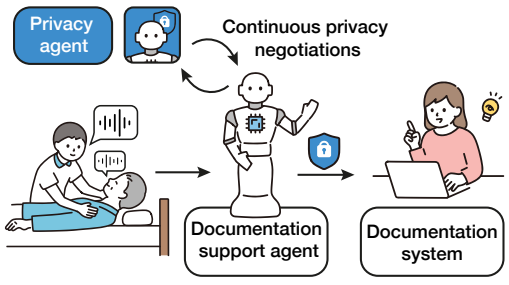

Figure 1: Introduction of a privacy agent for continuous privacy negotiations with the documentation support agent.[**]

Due to the sensitive nature of the application, we believe it is of utmost necessity to implement privacy protection measures. As also pointed out by other scholars (e.g., [18, 24]), we share the opinion that value-based design should be proactively pursued by

[**] Vectors and icons by Soco St in CC Attribution License via SVG Repo.

technology designers. As shown in Figure 1, we therefore want to add a privacy protection agent to the system. Instead of directly forwarding the transcript from the speech to the LLM, this additional software agent is consulted. The contained information should only be kept for the summary if none of the contextual privacy norms are violated. We propose a LLM-based solution, which requires two crucial steps: analyzing the information flows according to the key parameters of contextual integrity of spoken text (see Section 3.2) and taking into account the contextual privacy norms (see Section 3.3).

3.2 Analyzing Information Flows

The following is a transcript of one of the care scenarios conducted in the study described in Section 3.1. Significantly, our system does not currently facilitate speaker diarization between the care worker and the care recipient. In the present stage of the transcription model, we observed that diarization did not perform as effectively as anticipated.

> Good morning, Mrs. N. Good morning. Good morning. Morning, Mrs. N. And? Did you sleep well? Mrs. N. No. No, I didn't sleep well. Mrs. N., I'd like to get you dressed now because your daughter is coming to visit. Hm? Is that all right with you? Mrs. N., I'm going to... So, Mrs. N., I'm going to put you on now so that you speak nicer and everything. Then open your eyes. I'm in pain. Pain. Where is the pain, Mrs. Anna? There. Ah, not there. Does your shoulder hurt? Does the shoulder... Low back pain. I'll push you down a little more. There, slowly. That's better now, Mrs... Where's Gabi? Gabi is the daughter, right? [...]

With this brief excerpt from the transcript we demonstrate the feasibility of identifying some of the key parameters within information flows. The dataset `pii-masking-200k` [1], is designed to be used to fine-tune an LLMs to mask personally identifiable information (PII). This is a useful feature, e.g., to prevent users of an online forum to post telephone numbers publicly. In total, the dataset comprises 54 different PII classes, which constitute different types of sensitive data. Among them are *gender*, *sex*, *phone numbers*, *account numbers*, *IBAN*, *zip codes*, *passwords* and *social security numbers* that can be identified and masked by using this dataset [1]. We provided the excerpt from the transcript shown in Section 3.2 to an openly available model that has been fine-tuned on the English subset of the dataset, called `distilbert_finetuned_ai4privacy_v2` [19]. In Table 1 the results are shown. The LLM was able to successfully recognize the majority of the contained PII. Only the first occurrence of the prefix "Mrs." was not shown in the response.

Entity Group	Word	Position in text
PREFIX	mrs.	28–23, 73–77, 143–146, 258–262, 287–291, 423–427
FIRSTNAME	anna	428–432
	gabi	600–604, 606–610

Table 1: Identified PII from the excerpt of the transcript of the care documentation use case.

The parameters identified by such models can already be used for the analysis of an information flow (see Figure 2) to some extent. In practical terms, in relation to our use case, if the care recipient reveals sensitive details such as passwords or social security numbers, this information can already be excluded since it is not pertinent to care documentation. Assuming advancements in speaker diarization capabilities in future speech-to-text transformer models, it will become increasingly feasible to accurately identify the *sender* (i.e, the care worker and the care recipient), the *recipient* (i.e., the care worker who is using the recorded care action for the documentation), and the *subject* (i.e., the care recipient) by scanning for first names, last names, and prefixes. The *transmission principle* (i.e., confidential between care worker and care recipient) in the use case of care documentation support is a constant parameter, as long as no other information recipients or senders enter the context to whom the confidentiality does not apply, such as bystanders. In order to identify a complete information flow, algorithms must be used to analyze the links of the relevant parameters extracted by the LLM. Nevertheless, there are limitations, as the PII dataset aids to detect many, though not all types of attributes. While the dataset encompasses a wide range of PII classes, the identification of other *attributes* (e.g., some medication or vitals) necessitates the inclusion of these parameters in the dataset. Thus, in order to analyze all parameters of an information flow, PII is just a starting point.

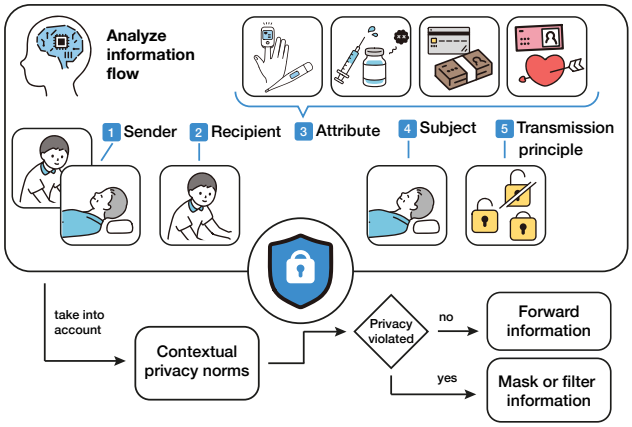

Figure 2: Flowchart of the privacy agent to provide privacy aware information.[**]

There are approaches that use transformer models to classify whether text contains sensitive topics of the categories of *politics*, *religion*, *health*, and *sexual habits* [29]. This is helpful to recognize whether sensitive data is affected in addition to the PII, but can not explicitly be used to recognize and mask e.g., vitals from a text, as whole sentences are classified rather than key terms. Moreover, a detailed differentiation of the affected attributes within the classification categories is desirable in order to make decisions for contextual integrity, as e.g., not all health related topics are subject to the same privacy norms. To our knowledge, no dataset currently exists that comprises an extensive collection of privacy-sensitive cases, and is annotated with the crucial parameters of the contextual integrity information flow. Given such a dataset, LLMs could

be fine-tuned to specifically analyze information flows, to recognize privacy-sensitive topics in conversations – a research focus worth exploring on the path to achieving privacy-aware robotics.

3.3 Contextual Privacy Norms

As can be seen in Figure 2, after analyzing the information flow it is necessary to take the applicable privacy norms into account. In the workshop setting described in Section 3.1, we also discussed privacy aspects. Attitudes towards privacy were mixed, in line with the findings outlined in Section 2.1. We identified that participants valued safety over privacy concerns in the care context. This sentiment is also shared among care workers, who view the protection of the residents' health as the highest value. Documentation also plays an important role in this context. Not only is it required by law for care workers, but it is also an imperative procedure for institutional control. While our proposed documentation support agent can improve documentation quality, recording audio during care routines poses a challenge to privacy that did not exist before. With the introduction of the privacy agent, we follow a value-based research direction, where we aim to achieve both goals: better documentation quality while maintaining privacy. Therefore, the privacy norms of the context have to be analyzed in detail. For example, in the context of care, it is acceptable for vital signs such as blood sugar or body temperature to be passed on, as the recipients are the care workers and this is relevant information. Nevertheless, if a care worker talks about private information of their life, such as family status, this is information which should not be included. The assessment based on the applicable contextual privacy norms determines whether the information is passed on, or filtered or masked (see Figure 2). The *transmission principle* is dependent on the context and the involved individuals, e.g., an information flow can be confidential between care workers and recipients, but not confidential if bystanders enter the scene.

Research has been conducted on privacy norms in certain contexts, e.g., Apthorpe et al. [5] investigate whether privacy norms of parents align with the regulation on Internet-connected "smart" children's toys. Other studies explore the privacy expectations of many-to-many human-robot interactions of autonomous vehicles and bystanders [8]. In order for the privacy agent to make a decision, the contextual norms have to be known. Therefore, as of now the privacy agent is limited to the predefined contextual norms of its use case. Nevertheless, if multiple contexts are researched and the privacy norms formulated, users of the robotic technology could activate the respective set of norms similar to a "Do not disturb" mode on a phone.

3.4 General Applications

The application of the conceptual privacy agent is not limited to the care documentation support use case. Given the contextual privacy norms are known where a robot operates, the privacy agent can be applied to other use cases and fields of HRI as well. With continuous real-time speech-to-text conversion (see Figure 3), the agent can be a middle man to check whether captured conversations constitute information flows according to the contextual integrity framework, and decide accordingly in real time to act. Compared to the use case outlined in Section 3.1, where it is possible to use the whole transcript for identifying information flows, in a real-time scenario words are continuously added to the transcript, and the privacy agent needs to continuously check chunks of the transcript, which requires significantly more computing power. Nevertheless, if realized, robots could react in real-time to privacy-sensitive situations, e.g., not forwarding information, turning off cameras, or moving out of a room.

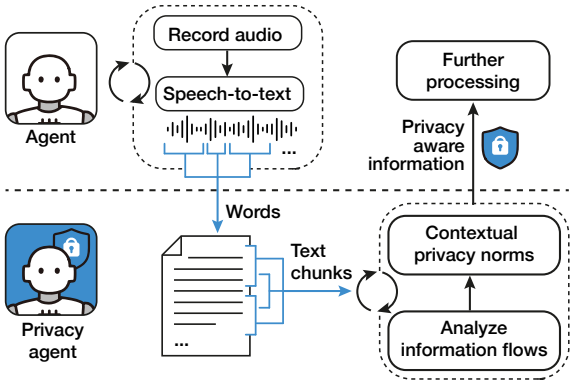

Figure 3: Conceptual real-time privacy negotiation between an agent and the privacy agent.[**]

4 DISCUSSION

We propose a conceptual *privacy agent* to detect whether an information flow is appropriate in the given context, keeping contextual integrity. However, as we rely solely on text-based LLMs, it is not possible to automatically identify a specific context. Thus, in the current state, the contextual norms must be known in advance and activated manually by the users of the proposed privacy agent. Only by combining several robotic abilities will the privacy agent be able to recognize a context, e.g., using biometric face recognition to know whether only family members are present. Although our contribution does not fully resolve the issue of privacy violations associated with the introduction of robotics, we believe it represents a promising direction. Robots capable of real-time recognition of privacy-sensitive content based on verbal communication present an opportunity to address various aspects of privacy. For instance, a robot programmed to deactivate sensors or exit the room during privacy-sensitive scenarios could mitigate concerns related to physical privacy.

ACKNOWLEDGMENTS

This research is supported by the Austrian Science Foundation (FWF) project Caring Robots // Robotic Care (CM 100-N). Vectors and icons by Soco St in CC Attribution License via SVG Repo.

REFERENCES

[1] ai4Privacy. 2023. Pii-Masking-200k (Revision 1d4c0a1). https://doi.org/10.57967/hf/1532
[2] Ahmad Alaiad and Lina Zhou. 2014. The Determinants of Home Healthcare Robots Adoption: An Empirical Investigation. *International journal of medical informatics* 83, 11 (2014), 825–840.
[3] Irwin Altman. 1975. The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding. (1975).

[4] 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 (2018), 1–23.

[5] Noah Apthorpe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents'{IoT} Toy Privacy Norms Versus {COPPA}. In *28th USENIX Security Symposium (USENIX Security 19)*. 123–140.

[6] Christopher Armbrust, Syed Atif Mehdi, Max Reichardt, Jan Koch, and Karsten Berns. 2011. Using an Autonomous Robot to Maintain Privacy in Assistive Environments. *Security and Communication Networks* 4, 11 (2011), 1275–1293.

[7] Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. 2023. WhisperX: Time-accurate Speech Transcription of Long-Form Audio. *INTERSPEECH 2023* (2023).

[8] Cara Bloom and Josiah Emery. 2022. Privacy Expectations for Human-Autonomous Vehicle Interactions. *31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)* (2022).

[9] Louis Brandeis and Samuel Warren. 1890. The Right to Privacy. *Harvard law review* 4, 5 (1890), 193–220.

[10] Judee K. Burgoon. 1982. Privacy and Communication. *Annals of the International Communication Association* 6, 1 (1982), 206–249.

[11] M Ryan Calo. 2011. Robots and Privacy. *Robot Ethics: The Ethical and Social Implications of Robotics* (2011), 187.

[12] George Demiris, Debra Parker Oliver, Jarod Giger, Marjorie Skubic, and Marilyn Rantz. 2009. Older Adults' Privacy Considerations for Vision Based Recognition Methods of Eldercare Applications. *Technology and Health Care* 17, 1 (2009), 41–48.

[13] Heather Draper and Tom Sorell. 2017. Ethical Values and Social Care Robots for Older People: An International Qualitative Study. *Ethics and Information Technology* 19, 1 (2017), 49.

[14] 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, 727–732.

[15] Álvaro García-Soler, David Facal, Unai Díaz-Orueta, Lucia Pigini, Lorenzo Blasi, and Renxi Qiu. 2018. Inclusion of Service Robots in the Daily Lives of Frail Older Users: A Step-by-Step Definition Procedure on Users' Requirements. *Archives of gerontology and geriatrics* 74 (2018), 191–196.

[16] Jan Philipp Harries. 2023. EM German (V01).

[17] Johan F Hoorn, Elly A Konijn, Desmond M Germans, Sander Burger, and Annemiek Munneke. 2015. The In-between Machine. In *Proceedings of the International Conference on Agents and Artificial Intelligence-Volume 2*. 464–469.

[18] Marcello Ienca, Tenzin Wangmo, Fabrice Jotterand, Reto W Kressig, and Bernice Elger. 2018. Ethical Design of Intelligent Assistive Technologies for Dementia: A Descriptive Review. *Science and engineering ethics* 24, 4 (2018), 1035–1055.

[19] Isotonic. 2024. Distilbert_finetuned_ai4privacy_v2.

[20] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825* (2023).

[21] Margot E. Kaminski. 2014. Robots in the Home: What Will We Have Agreed To. *Idaho L. Rev.* 51 (2014), 661.

[22] Helena Leino-Kilpi, Maritta Välimäki, Theo Dassen, Maria Gasull, Chryssoula Lemonidou, Anne Scott, and Marianne Arndt. 2001. Privacy: A Review of the Literature. *International journal of nursing studies* 38, 6 (2001), 663–671.

[23] 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 (2019), 412–434.

[24] Christoph Lutz and Aurelia Tamò. 2015. RoboCode-Ethicists: Privacy-friendly Robots, an Ethical Responsibility of Engineers?. In *Proceedings of the ACM Web Science Conference*. 1–12.

[25] Christoph Lutz and Aurelia Tamò. 2016. Privacy and Healthcare Robots–an Ant Analysis. *We Robot* (2016).

[26] Christoph Lutz and Aurelia Tamó-Larrieux. 2020. The Robot Privacy Paradox: Understanding How Privacy Concerns Shape Intentions to Use Social Robots. (2020).

[27] Helen Nissenbaum. 2004. Privacy as Contextual Integrity. *Wash. L. Rev.* 79 (2004), 119.

[28] Helen Nissenbaum. 2011. A Contextual Approach to Privacy Online. *Daedalus* 140, 4 (2011), 32–48.

[29] Michael Petrolini, Stefano Cagnoni, and Monica Mordonini. 2022. Automatic Detection of Sensitive Data Using Transformer-Based Classifiers. *Future Internet* 14, 8 (2022), 228.

[30] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*. PMLR, 28492–28518.

[31] Donald Redfoot, Lynn Feinberg, and Ari N Houser. 2013. *The Aging of the Baby Boom and the Growing Care Gap: A Look at Future Declines in the Availability of Family Caregivers*. AARP Public Policy Institute Washington, DC.

[32] Delphine Reinhardt, Monisha Khurana, and Luca Hernández Acosta. 2021. "I Still Need My Privacy": Exploring the Level of Comfort and Privacy Preferences of German-speaking Older Adults in the Case of Mobile Assistant Robots. *Pervasive and Mobile Computing* 74 (2021), 101397.

[33] Matthew Rueben, Alexander Mois Aroyo, Christoph Lutz, Johannes Schmölz, Pieter Van Cleynenbreugel, Andrea Corti, Siddharth Agrawal, and William D Smart. 2018. Themes and Research Directions in Privacy-Sensitive Robotics. In *2018 IEEE Workshop on Advanced Robotics and Its Social Impacts (ARSO)*. IEEE, 77–84.

[34] Lindsay H Ryan, Jacqui Smith, Toni C Antonucci, and James S Jackson. 2012. Cohort Differences in the Availability of Informal Caregivers: Are the Boomers at Risk? *The Gerontologist* 52, 2 (2012), 177–188.

[35] Elaine Sedenberg, John Chuang, and Deirdre Mulligan. 2016. Designing Commercial Therapeutic Robots for Privacy Preserving Systems and Ethical Research Practices within the Home. *International Journal of Social Robotics* 8, 4 (2016), 575–587.

[36] Amanda Sharkey and Noel Sharkey. 2012. Granny and the Robots: Ethical Issues in Robot Care for the Elderly. *Ethics and information technology* 14, 1 (2012), 27–40.

[37] Yan Shvartzshnaider, Noah Apthorpe, Nick Feamster, and Helen Nissenbaum. 2019. Going against the (Appropriate) Flow: A Contextual Integrity Approach to Privacy Policy Analysis. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 162–170.

[38] Daseul Yang, Yu-Jung Chae, Doogon Kim, Yoonseob Lim, Dong Hwan Kim, ChangHwan Kim, Sung-Kee Park, and Changjoo Nam. 2021. Effects of Social Behaviors of Robots in Privacy-Sensitive Situations. *International Journal of Social Robotics* (2021), 1–14.

[39] Salifu Yusif, Jeffrey Soar, and Abdul Hafeez-Baig. 2016. Older People, Assistive Technologies, and the Barriers to Adoption: A Systematic Review. *International journal of medical informatics* 94 (2016), 112–116.

Received 23 February 2024