Towards Designing Privacy-Compliant Social Robots for Use in Private Households: A Use Case Based Identification of Privacy Implications and Potential Technical Measures for Mitigation*

Björn Horstmann[a,b], Niels Diekmann[b], Hendrik Buschmeier[a,c] and Teena Hassan[c]

***Abstract*— Social robots are expected to increasingly appear in private households. The deployment of social robots in the private spheres of humans raises concerns regarding privacy protection. This paper analyses some of the legal implications of using social robots in private households on the basis of four practical use cases. It identifies the privacy concerns associated with each use case and proposes potential technical measures in the form of an initial concept for a companion privacy-app that could resolve or mitigate these concerns, and thereby enhance privacy compliance. The proposed app concept was evaluated in an exploratory study with ten participants. The preliminary results are encouraging and show that this concept has the potential to support the maintenance of privacy and provide control over the user's personal data and the robot's functions.**

I. INTRODUCTION

Hegel et al. [1] defined a social robot as "a robot plus a social interface", i.e. a robot possessing "social attributes by which an observer judges the robot as a social interaction partner." It possesses a "social form", serves a "social function", and is developed for a specific "social context" [1]. A few examples of social robots would include Pepper and NAO (both from SoftBank Robotics Group Corp., Tokyo, Japan), Jibo (Jibo, Inc., Boston, USA), and AIBO (Sony Corporation, Tokyo, Japan). Pepper and NAO have been used in various applications, ranging from healthcare to education. Jibo was meant to be an assistant at home, whose tasks were supposed to be e.g. weather forecasting and flight status checking. The pet robot AIBO resembles a dog and mimics its behaviour. Jibo and AIBO were commercially developed for use in private households to assist, accompany and entertain human beings in their daily lives.

To engage in social interactions with a human, a social robot should be capable of perceiving, learning, adapting to the human and its environment, producing gestures, expressing emotions, and communicating via natural language, among other things [2]. Perception and learning involve recording of data using different sensors and processing of recorded data using different methods. For example, in order to recognise a human user, a robot should be built with cameras and microphones as well as software for automatic face recognition and/or voice and speech recognition. In order to navigate through a household or personal space, a social robot should be able to map these spaces. In order to 'understand' a user's emotions, desires and intentions, and to respond accordingly, the robot should have the capability to model users' affective and cognitive states. This would require the robot to record, store and process information about the user and the interaction context. It is clear that the data that are recorded and processed by social robots deployed in private households pertain to humans and their personal living spaces. Therefore, the recording and processing of data can have serious legal implications, especially with regard to privacy and personal data protection.

In this paper, we discuss some of these legal implications on the basis of a set of use cases involving humans and social robots in private households. We suggest potential technical measures for mitigating these legal implications. The focus is laid on the development of a mock-up for a companion app (for a smartphone or tablet) that would support a social robot to function in greater compliance with the laws of privacy and personal data protection, and empower users to exercise greater control over the collection and processing of personal data and protection of their personal space. The keyword is 'privacy by design' (Art. 25 of the EU General Data Protection Regulation (GDPR)). In essence, it refers to the implementation of Technical and Organisational Measures (TOMs) – starting at the earliest stages of the design of the processing operations – which meet the principles of privacy and data protection (see Recital 78 GDPR).

Section II describes some of the relevant legal provisions that form the basis for privacy and data protection. Section III discusses the related work on privacy and data protection in the context of social robots. Section IV introduces the proposed app mock-up. Section V defines four use cases, identifies legal implications for privacy and data protection, and proposes privacy-by-design solutions. Section VI presents the results of an exploratory study evaluating the proposed app mock-up. Section VII concludes the paper.

II. LEGAL FRAMEWORK

The term 'privacy' encapsulates different concepts [3]. We do not examine any sociological or communicative approaches to privacy theory such as Nissenbaum's "contextual integrity" [4]. Instead, we focus on the legal perspectives on 'privacy', especially in the context of personal data protection. Privacy and data protection are related, but not identical.

*This research was supported by the German Federal Ministry of Education and Research (BMBF) in the project 'VIVA' (FKZ 16SV7959).

[a] Faculty of Linguistics and Literary Studies, Bielefeld University, Bielefeld, Germany, {bhorstmann, hbuschme}@uni-bielefeld.de

[b] Bielefeld University of Applied Sciences, Bielefeld, Germany, {bjoern.horstmann, niels.diekmann}@fh-bielefeld.de

[c] Social Cognitive Systems Group, Faculty of Technology, Bielefeld University, Bielefeld, Germany, {hbuschme, thassan}@techfak.uni-bielefeld.de

Data protection regulations are exclusively meant to govern the processing of personal data, which can evidently have privacy implications. However, not any aspect of privacy protection is covered under data protection law. The main sources of data protection law applied in this paper are international, supranational and constitutional law, such as Art. 8 of the Charter of Fundamental Rights of the European Union ("Protection of personal data"). It states:

1) "Everyone has the right to the protection of personal data concerning him or her."
2) "Such data must be processed fairly for specified purposes and on the basis of the consent of the person concerned or some other legitimate basis laid down by law. Everyone has the right of access to data which has been collected concerning him or her, and the right to have it rectified."
3) "Compliance with these rules shall be subject to control by an independent authority."

This Art. 8 has been concretised into a comprehensive body of law, called the GDPR (Regulation 2016/679/EU). The coming-into-force of the GDPR reflects a general development in law and has been replicated in several pieces of legislature around the world (for example, in the California Consumer Privacy Act (CCPA) or in the Biometric Information Privacy Act of Illinois (BIPA)). The focus of interest in this paper is directed to the GDPR. The GDPR is meant to govern the legal relation between controllers, processors, data subjects of personal data processing, and third-parties. The material scope of the GDPR mainly includes "the processing of personal data wholly or partly by automated means" (Art. 2 (1) GDPR). The GDPR lays out a number of principles and provides certain rights to data subjects. Art. 5 (1) of the GDPR states that "personal data shall be: (a) processed lawfully, fairly and in a transparent manner in relation to the data subject ('lawfulness, fairness and transparency'); (b) collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes (...) ('purpose limitation'); (c) adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed ('data minimisation'); (d) accurate and, where necessary, kept up to date" – "every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay ('accuracy'); (e) kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed (...) ('storage limitation'); (f) processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures ('integrity and confidentiality')."

As in any aspect of the law, there is rarely compliance without at least some kind of enforcement. Therefore, the GDPR grants the data subject the right to appeal against alleged unlawful processing of his or her data in the courts or in independent agencies (Arts. 72–79 GDPR). The unlawful processing may obligate the controller or processor to pay reparation to the appealing data subject, in case the unlawful processing has led to damages (Art. 82 GDPR). The processing of personal data shall be transparent at all times. Notwithstanding, the affected data subject has the right of access to his/her personal data and to be informed of the "purposes of the processing"; "the categories of personal data" that are processed; "the recipients or categories of recipient to whom the personal data have been or will be disclosed"; "where possible, the envisaged period for which the personal data will be stored, or, if not possible, the criteria used to determine that period"; and so forth (Arts. 13–15 GDPR). These legal principles and data subject rights are interwoven (see Table I). The controller is bound to uphold the legal principles, irrespective of the action of the data subjects. This includes the enactment of TOMs for personal data protection. This obligation does not directly apply to the developer or designer of a product, such as a social robot. However, the GDPR states that, "when developing, designing, selecting and using (...) products that are based on the processing of personal data or process personal data to fulfil their task, producers of the products (...) should be encouraged to take into account the right to data protection (...) with due regard to the state of the art, to make sure that controllers and processors are able to fulfil their data protection obligations" (Recital 78 S. 4. GDPR). Therefore, the general obligation to enact TOMs to safeguard data protection also applies to designers and producers, albeit indirectly. In the context of autonomous technology such as the aforementioned social robot, the line between producer and controller of data processing can be blurry. One could even envisage a joint responsibility of producer and controller [5]. In this paper, we examine how TOMs could be implemented during the design and development of a social robot in order to meet the data protection requirements laid out in Art. 25 of the GDPR that is concerned with data protection by design and by default.

TABLE I
SELECTED PRINCIPLES OF PERSONAL DATA PROTECTION AND SELECTED RIGHTS OF DATA SUBJECTS SPECIFIED IN GDPR.

Principles Binding Controllers	**Rights of Data Subjects**
Lawfulness, fairness	Rights to appeal, object, erasure
Transparency	Right of access, right to be informed
Purpose limitation, data minimisation, storage limitation	Right to restrict processing, right to erasure
Accuracy	Right to rectification

III. RELATED WORK

Rueben et al. [3] identified seven research themes for "privacy-sensitive robotics": data privacy, manipulation and deception, trust, blame and transparency, legal issues, domains with special privacy concerns, and privacy theory.

For these themes, various research directions were proposed. In this paper, we aim to address – with the help of four use cases – some of the privacy and legal issues identified by Rueben et al. [3]. Conveniently, most of these aspects are also covered by GDPR regulations. However, there is a crucial difference. Although, building a privacy-sensitive robot would be ethically desirable, it does not guarantee that it would be compliant with the law, by default. Therefore, it is important to address these issues within a legal framework, such as the GDPR.

Fosch-Villaronga et al. [6] summarise the results of discussions on ethical, legal and societal implications of social robots that took place at four international workshops on those topics during the period between 2015 and 2017. At these workshops, concerns about the "ability to control the collected and processed data", especially regarding data integrity, and "the ability to correct or amend the data", were raised. Discussions came to the conclusion that users would likely want to be able to "indicate which data can be collected" and that the collected data should be anonymised immediately, but that it could nevertheless be difficult for individual users to understand how their data are being processed and possibly even disseminated. The article names elements of the GDPR, such as information duties, individual rights and privacy by design, alongside technical and organisational measures as possible remedies to these problems, and specifically recommends making transparent which data are collected and giving users easy access to data and the ability to erase data. In this paper, we follow these recommendations and propose a concept for an app-based user interface to address the above-mentioned concerns when using social robots in private households.

Companion apps, similar to the one described in this paper, also exist for the commercially available social robot AIBO and for the voice assistant system Alexa (which provides similar functions to a social robot, albeit disembodied). Concerning transparency and management of personal data, the functionality of these companion apps is rather limited. They do not provide fine-grained control over data collected by the devices. The AIBO companion app allows users to configure whether their robot is allowed to take photos and upload them to a server and to browse or delete these photos. The Alexa companion app is restricted to giving users access to the voice commands that were uttered, deleting them, and providing feedback on misrecognised utterances in order to improve the service. The concept of an app-based user interface for a social robot presented in this paper goes beyond such basic functionality.

IV. PROPOSED APP-BASED USER INTERFACE

This section introduces the proposed mock-up for an app-based user interface for enhancing the privacy compliance of social robots. It was conceived, designed and developed to enable users to have greater control over their personal data by informing them about what data are being processed and which new information has been learned by a social robot that they are using in their private households. A robot's companion app following this mock-up can be seen as a technical measure for implementing the requirements of privacy by design (Art. 25 GDPR). When implemented, it will serve as an additional software that interfaces and communicates with the software running on and controlling the social robot, in order to provide transparency and enhance accuracy in personal data processing.

The mock-up consists of 14 different screens, each designed to fulfil a different purpose. There are screens that record user consent and enable the user to control the privacy settings as well as view/edit/delete the different information learned by the robot about the user and other persons in their social circle. How the functionalities provided via these screens could contribute towards mitigating the data protection risks, will be discussed in depth in Section V. The proposed app mock-up was created using the low-fidelity wireframing software Balsamiq (Balsamiq Studios, LLC, Sacramento, USA), which supports the creation of simple user interfaces that are tailored to different mobile phones and tablets. In addition to the aforementioned functionalities related to privacy and data protection, there was a focus on usability based on selected human-system interaction principles from ISO 9241-110:2006. Furthermore, an analysis of functional and nonfunctional requirements was done.

Fig. 1A presents the start screen of the proposed app mock-up. It provides a high-level view of the different functionalities included in the app mock-up. As can be seen, there is a button called "Current view of my robot". Like in the AIBO app, by clicking on this button, users would be able to see through the 'eyes' of the robot. That is, users would be able to see what the robot's cameras are currently viewing/capturing via a small window that shows the live video stream. The next button is named "Current position of my robot". Its function is rather self-explanatory, and its role in privacy protection will be examined in detail in Section V (see Fig. 1C). The button "Last learned information" would take the user to a screen that provides an overview of the information that the robot has learned about the user so far (see Fig. 1B). This includes data categories like personal information, biometric data, hobbies, interests, social contacts and events. This screen would enable users to understand the data processing of the robot in a transparent and clear way. The button "Overview of processed personal data" would lead users to a screen that serves a similar purpose, but shows data that are used by the robot to fulfil its assigned functions. This includes visual data, auditory data, position data and conversational data. In addition to these, there are screens to configure the privacy settings and to explicitly record user consent. These will be described in the following section.

V. USE CASES & PRIVACY-BY-DESIGN SOLUTION

In this section, we present four use cases involving humans and social robots in a single person private household setting, and highlight various issues related to privacy and data protection that can arise in this context. The main questions that we considered are: (i) how might social robots violate a user's privacy, (ii) in which cases might the processing

of personal data become problematic, and (iii) what privacy concerns might users have – even if irrational or unprompted – about their privacy that this app could address? For each use case, we will discuss potential solutions to address the identified issues. These solutions are proposed as functionalities designed and integrated in the app mock-up introduced in Section IV as well as in the form of functions of the robot. These solutions are designed on the basis of the legal principles and rights of data subjects defined in the GDPR. The use cases, privacy concerns and proposed technical solutions are summarised in Table II.

A. Use Case 1 – First interaction

This use case relates to the very first interaction between the social robot and its owner. The robot gathers information in order to get to know the user (e.g. their names). It is assumed that the robot first learns to recognise its owner by capturing their facial features. This information will be stored by the robot so that it is able to use this personal data for future interactions with the user and to distinguish its owner from other persons. This initial interaction already entails various legal implications. The user's name qualifies as personal data, since personal data is defined in Art. 4 No. 1 GDPR as "any information relating to an identified or identifiable natural person." The collection of such data constitutes as "processing" under the meaning of Art. 4 No. 2 GDPR, because it is an "operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means." The processing of personal data shall be legal, only if and to the extent that at least one of the legal bases of Art. 6 GDPR applies. The sole legal basis that is applicable in this particular scenario is that the processing of personal data takes place on the grounds of the user's consent (Art. 6 (1) S. 1 lit. a)). The consent of the data subject refers to "any freely given, specific, informed and unambiguous indication of the data subject's wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her" (Art. 4 No. 11 GDPR).

Furthermore, the facial features collected and stored by the robot are not mere personal data, but belong to the special category of personal data, referred to as 'sensitive data', defined in Art. 9 GDPR. This includes, inter alia, data concerning health and biometric data that can be used to uniquely identify a natural person. Facial dimensions belong to such biometric data. The processing of such sensitive data upon the grounds of the data subject's consent shall only be allowed, if such "consent to the processing of those personal data for one or more specified purposes" has been given *explicitly* (Art. 9 (2) lit. a) GDPR). An implied consent is not legally sufficient. Therefore, the designer or producer of a social robot must implement necessary measures to obtain the consent of the data subject in a legally appropriate way. Furthermore, when designing a social robot, it should be taken into account that the "data subject shall have the right to withdraw his or her consent for processing personal data at any time" and that "it shall be as easy to withdraw as to give consent" (Art. 7 (3) S. 1, 4 GDPR).

Proposed solutions: Drawing inspiration from existing apps, a dialog box that appears on the app interface, when the robot is used for the first time by a new user, could be a possible solution to obtain explicit consent. Through this dialog box, users can accept or reject the 'Terms of Use' that are related to the usage of the social robot, which explicitly mentions the processing of personal data. Accepting the 'Terms of Use' allows the robot to process the user's data in the first place. This can be seen as the first requirement to be considered when starting the robot. The robot's functions will be limited if these terms and conditions are rejected. For instance, the processing of biometric data will be switched off by default. Still, these functions could be activated later on, if the terms and condition are accepted at a later stage via the app interface. Acceptance of the 'Terms of Use' by clicking the 'I Agree' button can be seen as an explicit consent given by the user. This ensures that the processing of sensitive data like facial information takes place legally. Thus, the privacy concerns and risks associated with Use Case 1 can be mitigated. Giving an explicit consent would also correspond to an "opt-in" regime, as proposed by Rueben et al. [3]. The robot won't collect or share personal information unless the user explicitly permits/opts for it.

B. Use Case 2 – Telephone conversation

The second use case relates to the scenario where the user is talking to his/her physician on the phone about some health problems. At the end of the phone call, the user makes an appointment to visit the doctor at a later time. The social robot captures these pieces of information related to the user's health and medical appointments via microphones and processes and saves the data. Personal data concerning health are sensitive data, and the recording of this data is legally equivalent to the processing of the data. In the absence of any other legal basis, it is mandatory to obtain the user's explicit consent in order to process this data. Furthermore, this explicit consent must be given voluntarily prior to the processing of the data. Consent given subsequently would not legalise the processing retroactively. In order to enable the processing of 'sensitive data' of this kind in such circumstances, or to avoid, or at least minimise, the collection of such data, developers are bound to implement technical measures accordingly.

Proposed solutions: Given that the user explicitly agreed to the terms of use of the robot which includes processing of personal data, the processing of the audio/speech data during telephone conversation can be perceived as legally appropriate. However, the user can still be provided additional means to proactively or reactively control the processing of such data. A proactive solution could be to provide the user an easy-to-use interface to deactivate the robot's microphone in advance, so that it will not capture audio/speech data during the subsequent phone call. Such an interface could be provided as part of the app, as shown in Fig. 1D. A reactive solution could be to provide suitable interfaces to allow users to view and delete stored conversation data

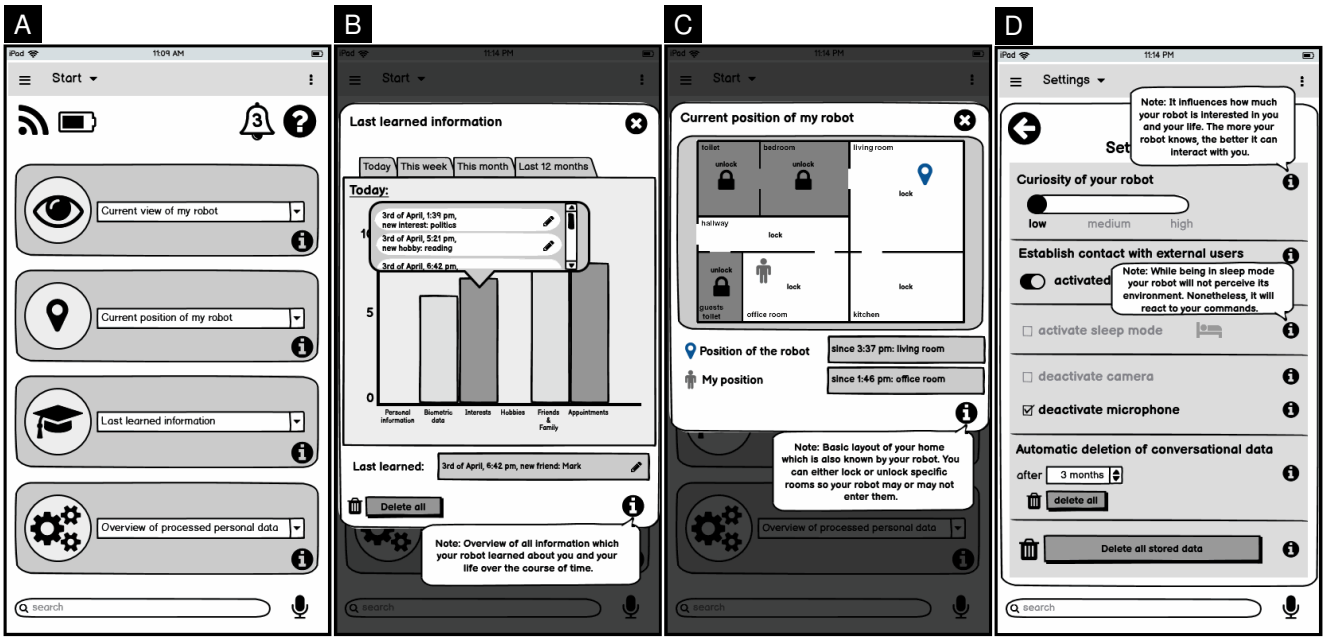

Fig. 1. Selected screens from the proposed app mock-up: A - Start screen; B - Last Learned Information screen; C - Current Position of Robot screen; D - Settings screen [Note: All 14 screens are available in higher resolution in supplementary material [7] and were originally in German.]

TABLE II

A SET OF FOUR CONCEIVABLE USE CASES, ASSOCIATED PRIVACY CONCERNS, AND POTENTIAL TECHNICAL MEASURES FOR MITIGATION.

#	Description	Privacy & data protection concerns	Proposed technical measures
1	First interaction: Robot meets the user for the first time.	– Lack of prior consent of the data subject. – Processing of sensitive biometric data.	– An app-based user interface to (i) show the terms of use which includes personal data processing, and (ii) to obtain explicit consent of the user.
2	User has a telephone conversation with physician.	– Potential lack of explicit consent. – The (explicit) consent given by the data subject previously is insufficient.	– An app-based user interface to allow the user to explicitly mute the robot's microphone and to delete all stored conversation data afterwards, as illustrated in Fig. 1D.
3	User gets a visitor.	– Lack of consent for data processing of a third person. – Lack of consent for data processing of a third person's biometric data.	– An app-based user interface to put the robot to sleep or to prohibit establishing contact with external users, as illustrated in Fig. 1D. – Identifying the robot owner via voice recognition to distinguish between data subjects.
4	Robot enters private spaces such as bathrooms unsolicited.	– Presumed violation of the data subject's privacy.	– An app-based user interface to allow the user to explicitly lock specific rooms in order to prevent access to robot, as illustrated in Fig. 1C. – A robot, whose curiosity level can be controlled by the user via the app-based interface, as illustrated in Fig. 1D, and who has a low curiosity level by default.

afterwards, so that they can remove any inappropriate or sensitive information that might have been captured by the robot. Furthermore, options can be provided to configure the robot such that it automatically deletes all conversation data after a specific period of time, which can be set individually by the user via the app interface.

C. Use Case 3 – Visitor

The third use case is divided into two different scenarios. In the first scenario, it is assumed that the robot owner is visited by a friend. The second scenario defines the visitor as a craftsperson. In both scenarios, personal data of visitors could be processed by the social robot (e.g. their speech and face information can be recorded). From a legal perspective, this use case has to be analysed with respect to the nature of the visits or specifics of the visitors. According to Art. 2 (2) lit. c), the GDPR "does not apply to the processing of personal data (...) by a natural person in the course of a purely personal or household activity" (household exemption). In other words, this is the case, if the visit is not connected to a professional or commercial activity. Therefore, assuming that the user is in control of the processing of personal data, the first scenario (being visited by a friend) would not fall under the material scope of the GDPR as stated in Art. 2 of the GDPR. In contrast to this, the household exemption does not apply to the visit of

a craftsperson in a strict or partially professional capacity. In the absence of another legal basis providing sufficient grounds for the processing of the craftsperson's personal data including sensitive personal data, the robot user must be given the means to obtain the explicit consent of such visitors or to avoid their data processing altogether.

Proposed solutions: Fig. 1D shows various settings that can be controlled by the user via the app interface. One of the options enables the user to prohibit the robot from establishing contact (e.g. greeting, starting conversations, etc.) with persons other than the robot owner/user, with the help of technology such as face and voice recognition (although not fool-proof). Consequently, personal data of persons other than the user will not be processed by the social robot. This might be a solution especially for the second scenario, so that the personal data of craftspersons are not processed without their consent. In order to ensure privacy by default, the option "Establish contact with external users" can be set by default to 'deactivated' and recommended as the standard setting.

Obviously, turning off the robot or putting it into the sleep mode could also be a solution in these scenarios. Enabling/Disabling the sleep mode is also possible via the 'Settings' screen in Fig. 1D. But, this would transfer the data controller's responsibility for lawful data processing to the data subject, which violates GDPR regulations. However, if the robot does not exchange personal data online and the visit is exclusively related to private activities, then conceivable risks and problems through data processing as stated above are nonexistent, due to the already mentioned household exemption. Furthermore, in order to avoid the recording of speech not spoken by the robot owner, voice recognition could be used to distinguish between different data subjects and to ensure as far as possible that only the speech of the robot owner is processed. However, in order to do this, explicit consent of the user would be necessary (similar to User Case 1). Preventing the processing of personal data that does not belong to the intended user complies with the requirement of the GDPR to minimise data processing to the bare minimum that is needed to fulfil the intended purpose. However, the capture of personal data of such data subjects might not be problematic as long as that information is deleted immediately by the social robot by default.

D. Use Case 4 – Entering private spaces

The fourth use case relates to the scenario where the social robot enters the robot owner's bathroom without asking for permission while the person (data subject) is changing clothes. In general, this use case is connected more with privacy violations than data protection regulations. The violation of privacy in this use case is clear, even if the theme is rather subjective. The right to privacy as the manifestation of the inherent right of personality [8] entails everybody's right to be left alone in his or her refugium, irrespective of specific rules. This right to be left alone in the confined space of one's home should be acknowledged in the context of product development. Therefore, a concept for self-restraint in accordance with the user's wishes should be implemented within the design of a social robot on principled grounds.

Proposed solutions: A possible solution for avoiding such privacy violations could be to provide an app interface, such as the one shown in Fig. 1C. Here, the app interface allows users to decide which rooms the social robot is allowed to enter. To realise this solution, the robot needs knowledge of the basic layout of the user's home so that it knows its current location (room) and which rooms are forbidden to enter. This information can then be used during path planning so that forbidden rooms are avoided during navigation. This information (robot's location, layout of rooms, forbidden rooms) is represented visually in the app interface shown in Fig. 1C. By clicking on the corresponding room, it can be 'locked' or 'unlocked' for the robot. A symbol in form of a black lock appears to indicate that the room is 'locked' or forbidden for the robot to enter. This also informs the user that the robot will stay outside the room. Hence, if all other rooms are 'locked' at the request of the user, the robot will not leave the current room in which it is located. A 'locked' room can be 'unlocked' for the robot by clicking on the lock symbol. After unlocking a room, the black lock symbol disappears, indicating that the social robot might enter that room. In this way, users would be able to adjust the privacy settings associated with the robot's navigation so that it fits their individual idea of privacy.

Additional control can be given to the user to control the privacy settings via the 'Settings' screen of the app interface shown in Fig. 1D. Here, the user can define the level of 'curiosity' of the robot, by setting it to low, medium or high. This can be seen as a way to enforce privacy protection and could correspond to "levels of clearance", as stated by Rueben et al. [3]. This setting influences how much the robot is interested in the user's life and to what extent it tries to learn about the user. For the purpose of ensuring privacy by default, the curiosity level of the robot is initially set to 'low'. Users can configure it later to match their idea of privacy. For now, it is not determined how the different 'curiosity levels' affect the robot's functions in detail. But it is conceivable that it would influence the depth of Human-Robot Interaction (HRI) in general, or more specifically, what kind of personal information will be processed by the robot. An exact definition and explanation of this feature is omitted on purpose so that the participants of the study could express their expectations regarding this feature.

VI. EVALUATION

The app mock-up presented in Sections IV and V was evaluated via a contextual interview lasting about 30 minutes and a short questionnaire. The interview facilitated a qualitative evaluation and the questionnaire enabled a quantitative evaluation of the functions of the proposed app mock-up. In total, ten participants (all of them students in Bielefeld; 5 female, 5 male; mean age: 21.9 years) consented to and partook anonymously in the evaluation. At the beginning, the participants were asked to imagine that they own a social robot that gives them company in their everyday life. In

addition to this, they were shown a picture of a social robot in a living room. The interviewer then told them that the robot can learn personal information about its user, and that they will soon see a concept for a companion app with which they can control the robot and its functions. After this briefing, they were shown print-outs of the relevant 12 screens (see [7]) from the mock-up. They examined these screens and during that time, they were encouraged to express their opinions by 'thinking aloud'. After this, the interviewer posed them questions related to the functions and design of the app mock-up, and asked them their opinion about using the robot at home. The questions included, for example, whether the functions of the user interface elements in the mock-up were clear, whether the participants would make use of certain functions provided in the app mock-up, etc. After the interview, the participants filled a questionnaire.

During the interview, six participants stated that some screens look familiar and resemble other apps they know. Others mentioned that they would have preferred additional dialog boxes in order to confirm user input, e.g. when deleting data. However, in general, the participants expressed concerns regarding the storage of (personal) data by the robot. Most of such concerns were due to media reports about scandalous handling of personal data, for example, by social media service providers. In general, data storage in a 'cloud' was not considered safe, and two participants mentioned that robots with internet connection would worry them. Local data storage on the robot itself was not considered to be too problematic, but reliability of the data controller was a major concern among the participants. Participants were familiar with the potential risks associated with the disclosure of personal data when using web services, but their willingness to disclose personal data generally varied from person-to-person. Participants stated that they would reveal personal information such as hobbies and interests, as long as they felt that the data are stored securely and are necessary for the robot to fulfil its functions. Three participants stated that they would use the function "Curiosity of your robot" depending on their day; for instance, if they have visitors at home or if they feel like talking to someone. Regarding this, one participant stated (translated from German): *"I find it quite nice that you can somewhat adjust how strong the robot is interested in your life. It is stupid, if it sticks to your heel all the time (...)"*. All participants had an idea or expectation about what this setting could do, even if it was not clear to some of them how it would affect the robot's behaviour. With regard to using social robots in private households, two participants expressed strong hesitation, especially in communicating and interacting with the robot. These participants also questioned the benefits of owning a social robot. The participants were also questioned about their understanding of the term 'privacy'. The term was often used to refer to their own flat or specifically to particular rooms, like bathroom or bedroom. It was also associated with aspects of private life, such as intimacy with the partner, changing of clothes, and very intimate conversations between friends or about certain health-related appointments.

TABLE III

RELEVANT ITEMS SELECTED FROM THE QUESTIONNAIRE USED FOR EVALUATING THE APP MOCK-UP IN THE EXPLORATORY STUDY.

#	**Item** (translated from German)
8	It does not bother me that the robot knows so many personal things about me.
9	The robot would restrict my privacy.
10	I think it's good if the robot captures and saves information about me on its own.
11	Sensitive data (such as account numbers, passwords, religious beliefs, etc.) should be saved by the robot.
12	The shown app helps me to maintain my privacy.
13	My data are clearly displayed by the app.
14	Data processing and storage happens transparently.
15	I have full control over my data.
17	All app functions are understandable.
19	The app makes me feel that I have control over the robot and its functions.

Following the interview, the participants were asked to fill out a questionnaire which covered the topics: (i) usage of the robot, (ii) privacy by design in the sense of transparency, (iii) control over personal data, and (iv) how the app concept supports these features. The questionnaire consisted of a set of 22 pre-defined statements (see [7]), of which 14 had to be rated on 5-point Likert-scales. Participants had to choose whether they were undecided (0), or whether they agreed (2), somewhat agreed (1), disagreed (-2) or somewhat disagreed (-1) with a particular statement. The scores of an answer were averaged over all participants in order to obtain an indication of the overall degree of approval. Table III shows a selected set of 10 statements related to privacy protection, control over personal data, and functions of the app mock-up.

The results of items 8, 9 and 10 indicated a mixed response regarding privacy protection and social robot usage (mean score: -0.2, 0.2 and 0.4, respectively). Item 11 showed a disapproval of the storage of sensitive data by the robot (mean: -1.5). All the other items are related to the functions and design of the app mock-up and showed higher levels of approval among participants. For instance, all participants either agreed or somewhat agreed to item 13 (mean: 1.7), and therefore they most likely found the app mock-up and its interface elements to be clear. An exception is represented by item 14. The result for this item (mean: 0.6) indicated that the participants desire more transparency in data processing and storage via the app. In contrast, the participants somewhat agreed that they had control over the stored data via the app (item 15, mean: 1.1). One of our hypotheses was that the maintenance of privacy and user's control over the robot and its functions via the app are positively correlated. To test this hypothesis, Pearson's correlation coefficient r was computed between items 12 and 19. A value $r = 0.53$ was obtained, showing a moderate positive linear correlation between the two items. This preliminary result is encouraging, since it shows that the user having control over the robot's functions via the proposed app could help in maintaining privacy.

To summarise, the interview revealed that the overall attitudes of the participants towards self-learning social ro-

bots were rather mixed. For instance, interview statements (translated from German) ranged from *"Yes, firstly, I find it to be very interesting. It is totally cool, if you can see all the information, which it finds out about you."* to *"So, does it move? A little bit creepy in general. (...) If I know what has been recorded and how it was linked and processed, it makes things less creepy (...)"*. This was also reflected in the mean scores for the questionnaire items 8, 9 and 10 (see paragraph above). The proposed app mock-up was received positively with respect to clarity, maintenance of privacy, control over data, and control over robot's functions (mean scores above 1.0 for items 12, 13, 15, 17 and 19). More qualitative statements from participants regarding the screens in Fig. 1 are provided in the supplementary material [7]. A more elaborate evaluation with a larger sample size representing different demographic cohorts is necessary for a more conclusive evaluation of the proposed app concept.

VII. CONCLUSION

In this paper, we analysed some of the legal implications (mainly based on GDPR) – regarding privacy and data protection – of using social robots in private households. We identified privacy concerns based on four concrete use cases that represent typical situations in the life of a person owning a social robot. We propose that one way to address these concerns with 'technical measures' would be to accompany social robots by a 'privacy app'. It could help make transparent which (types of) data the robot collects, stores and analyses, and, at the same time, could give owners detailed control over this data by providing intuitive means to inspect and erase (individual) data points or interaction episodes. The paper presents a concept for the user interface of such an app, and provides the results of an exploratory study on whether it could support the maintenance of privacy and provide control over users' data and the robot's functions.

The technical measures taken are an explication and concrete implementation of the general recommendations regarding privacy, data protection and social robots proposed in a recent series of international workshops [6]. The proposed app mock-up also realises the suggestions from Rueben et al. [3] for developing interfaces that enable robot users to specify their privacy preferences according to the situation. Although such detailed concerns about privacy are less relevant for social robots developed and/or deployed as part of research projects (see [9]), the concept presented here could serve as a useful blueprint for developers of commercially sold social robots. In order to meet certain GDPR regulations that cannot be solved by a companion app, technical and/or organisational measures need to be realised on the robots themselves. Clarifications and acquisition of consent, e.g. before processing face data, could for example be realised via dialogue based interactions – which implies that the robot itself should be given a sense of the privacy implications of its actions and how to mitigate them. A problem then arises about how to classify information so that the robot knows which information can be shared and with whom. Here, a social robot which can detect and distinguish certain social contexts within its software architecture might be needed, as Rueben et al. suggested [3].

We aim to implement and integrate these concepts within the architecture of a social robot [10], whose behaviour and actions will be (partially) driven by previous interaction episodes retrieved from memory [11], and which will be able to make its behaviours and their underlying causes transparent to its users by means of verbal explanations [12]. As the technical development progresses, further information and features will be added to the app concept. In this paper, the focus was limited to social robots in single person private households. Future work could extend this approach to assess risks in the context of multi-person households.

ACKNOWLEDGEMENTS

We thank Stefan Kopp, Head of Social Cognitive Systems Group, Bielefeld University, for providing the opportunity to work on this topic; Sonja Stange for sharing a few tips on writing; and Bielefeld University of Applied Sciences for the consultation on the legal aspects involved in this work.

REFERENCES

[1] F. Hegel, C. Muhl, B. Wrede, M. Hielscher-Fastabend, and G. Sagerer, "Understanding social robots," in *2009 Second International Conferences on Advances in Computer-Human Interactions*, 2009, pp. 169–174.

[2] T. Fong, I. Nourbakhsh, and K. Dautenhahn, "A survey of socially interactive robots," *Robotics and Autonomous Systems*, vol. 42, pp. 143–166, 2003.

[3] M. Rueben, A. M. Aroyo, C. Lutz, J. Schmölz, P. Van Cleynenbreugel, A. Corti, S. Agrawal, and W. D. Smart, "Themes and Research Directions in Privacy-Sensitive Robotics," in *2018 IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*, Genova, Italy, 2018, pp. 77–84.

[4] H. Nissenbaum, "Privacy as contextual integrity," *Washington Law Review*, vol. 79, no. 1, pp. 119–157, 2004.

[5] B. Wagner, "Disruption der Verantwortlichkeit. Private Nutzer als datenschutzrechtliche Verantwortliche im Internet of Things," *Zeitschrift für Datenschutz*, no. 7, pp. 307–312, 2018.

[6] E. Fosch-Villaronga, C. Lutz, and A. Tamò-Larrieux, "Gathering expert opinions for social robots' ethical, legal, and societal concerns: Findings from four international workshops," *International Journal of Social Robotics*, 2019.

[7] B. Horstmann, N. Diekmann, H. Buschmeier, and T. Hassan, "Supplementary material for the publication "Towards Designing Privacy-Compliant Social Robots for Use in Private Households: A Use Case Based Identification of Privacy Implications and Potential Technical Measures for Mitigation"," Bielefeld, Germany, 2020. [Online]. Available: https://doi.org/10.4119/unibi/2944720

[8] B. van der Sloot, "Privacy as personality right: Why the ECtHR's focus on ulterior interests might prove indispensable in the age of "Big Data"," *Utrecht Journal of International and European Law*, vol. 31, no. 80, pp. 25–50, 2015.

[9] E. Krempel and B. Steckler, Eds., *GUIDE. Leitlinien für den Datenschutz in der wissenschaftlichen Forschung zu Aspekten der Mensch-Technik-Interaktion*. Karlsruhe, Germany: Fraunhofer IOSB, 2019.

[10] S. Stange, H. Buschmeier, T. Hassan, C. Ritter, and S. Kopp, "Towards self-explaining social robots: Verbal explanation strategies for a needs-based architecture," in *Proceedings of the Workshop on Cognitive Architectures for HRI: Embodied Models of Situated Natural Language Interactions (MM-Cog)*, Montréal, Canada, 2019.

[11] T. Hassan and S. Kopp, "Towards an interaction-centered and dynamically constructed episodic memory for social robots," in *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, Cambridge, UK, 2020, pp. 233–235.

[12] S. Stange and S. Kopp, "Effects of a social robot's self-explanations on how humans understand and evaluate its behavior," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, Cambridge, UK, 2020.