Social Robot Co-Design Canvases: A Participatory Design Framework

MINJA AXELSSON, Department of Computer Science and Technology; University of Cambridge, UK,
Group on Artificial Intelligence for People and the Society (GAIPS), INESC-ID, Portugal
RAQUEL OLIVEIRA, Group on Artificial Intelligence for People and the Society (GAIPS), INESC-ID,
Iscte-University Institute of Lisbon (CIS-IUL), Portugal
MATTIA RACCA, Idiap Research Institute, Switzerland
VILLE KYRKI, Department of Electrical Engineering and Automation, Aalto University, Finland

Design teams of social robots are often multidisciplinary, due to the broad knowledge from different scientific domains needed to develop such complex technology. However, tools to facilitate multidisciplinary collaboration are scarce. We introduce a framework for the participatory design of social robots and corresponding canvas tool for participatory design. The canvases can be applied in different parts of the design process to facilitate collaboration between experts of different fields, as well as to incorporate prospective users of the robot into the design process. We investigate the usability of the proposed canvases with two social robot design case studies: a robot that played games online with teenage users and a librarian robot that guided users at a public library. We observe through participants' feedback that the canvases have the advantages of (1) providing structure, clarity, and a clear process to the design; (2) encouraging designers and users to share their viewpoints to progress toward a shared one; and (3) providing an educational and enjoyable design experience for the teams.

CCS Concepts: • **Human-centered computing** → **User interface toolkits**; **Participatory design**; *HCI theory, concepts and models*; *Interaction design theory, concepts and paradigms*; User centered design;

Additional Key Words and Phrases: Human-robot interaction, social robots, robot design, participatory design, design tool

ACM Reference format:

Minja Axelsson, Raquel Oliveira, Mattia Racca, and Ville Kyrki. 2021. Social Robot Co-Design Canvases: A Participatory Design Framework. *Trans. Hum.-Robot Interact.* 11, 1, Article 3 (October 2021), 39 pages. <https://doi.org/10.1145/3472225>

This project was funded by Futurice. M.A. acknowledges a grant from the project AMIGOS:PTDC/EEISII/7174/2014, and the Osk. Huttunen Foundation, and EP/T517847/1 (EPSRC). R.O. also acknowledges a grant by Fundação para a Ciência e a Tecnologia: (FCT) from the project AMIGOS:PTDC/EEISII/7174/2014 and a Ph.D. grant with the following reference: PD/BD/150570/2020. The work of M. R. and V. K. was partially financially supported by Academy of Finland Strategic Research Council project ROSE, 314180.

Authors' addresses: M. Axelsson, University of Cambridge, Department of Computer Science and Technology, The Computer Laboratory, William Gates Building, 15 JJ Thomson Avenue, Cambridge CB3 0FD; email: minja.axelsson@cl.cam.ac.uk; R. Oliveira, Group on Artificial Intelligence for People and the Society (GAIPS), INESC-ID, Iscte-University Institute of Lisbon (CIS-IUL), Av. das Forças Armadas, 1649-026, Lisbon, Portugal; M. Racca, Idiap Research Institute, Rue Marconi 19, CH - 1920, Martigny, Switzerland; email: mattia.racca@idiap.ch; V. Kyrki, Aalto University, Department of Electrical Engineering and Automation, Maarintie 8, 02130 Espoo, Finland; email: ville.kyrki@aalto.fi.

This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License.

© 2021 Copyright held by the owner/author(s).
2573-9522/2021/10-ART3 \$15.00
<https://doi.org/10.1145/3472225>

1 INTRODUCTION

Designing social robots requires a deep understanding of human behaviour and intelligence, as well as a diverse set of technical skills. Citing social robotics pioneers Cynthia Breazeal, Kerstin Dautenhahn, and Takayuki Kanda, “*the design of social robot technologies and methodologies are informed by robotics, artificial intelligence, psychology, neuroscience, human factors, design, anthropology, and more*” [17, p. 1935].

As a consequence of this multidisciplinary, managing the design process of diverse teams that are required to build such complex technology can be a difficult task.

When multidisciplinary teams work together, phenomena such as lack of clarity of team goals, as well as difficulty of sharing information and creating shared understanding due to the lack of a common terminology can be substantial challenges to the design process (e.g., References [33, 53, 91, 99]). Co-operation requires distributed cognition, where groups of minds interact with each other through tools and artifacts, with the goal of creating a shared product that reflects each individuals’ expertise [36]. Although approaches that address these problems have been proposed in the past (see Section 2), complete, easy-to-use tools or artifacts that facilitate multidisciplinary collaboration to design social robots are still scarce.

In this article, we attempt to address challenges present in multidisciplinary collaboration, by providing a tool that can be used by researchers and designers throughout the process of designing a social robot. The proposed tool is a design framework, detailing the parts and considerations of designing a social robot. The overall structure of the framework, pictured in Figure 1, can be separated into three phases: (1) the problem space, (2) the design guidelines, and (3) the solution space. To function as a collaborative tool, the framework is created into a corresponding tangible form as a series of canvases, constructed for co-design participants to work on together. The goal of the canvases is to facilitate the discussion of the design team around the task at hand, i.e., the design of a social robot. The framework intentionally goes beyond the technical implementation of the robot design to be able to involve non-technical stakeholders in the co-design process.

We argue that the proposed tool (1) improves the clarity of the design process, (2) enables designers and users to share their expertise and viewpoints with each other in a productive manner and reach a shared viewpoint, and (3) contributes toward an educational and enjoyable nature of the design process. Additional advantages of the framework are also discussed: being a flexible context tool (i.e., being applicable to multiple contexts where a robot can potentially operate), including ethical considerations into the design process, and being modifiable. The tool is made available for public use in Reference [4].

We evaluated the implementation of the proposed framework in two design processes. In two pilot studies, multidisciplinary teams designed a gaming robot for teenagers, and a robot for guiding customers at a library. An evaluation of the framework was conducted with 15 people who took part in the design processes, including roboticists, domain experts, and prospective users of the robots. The evaluation revealed several advantages related to the use of our framework. In particular, the majority of the participants indicated benefits related to the structure and clarity of the design process, by indicating that the use of the framework increased their ability to think clearly, their ability to conduct an exhaustive review on the subject of social robot design and that the framework provided a clear, easy-to-follow structure. The participants also indicated that the use of the framework made the experience of designing the robots more enjoyable and educational, putting an emphasis on the versatile ability to share different viewpoints and engage in rich multidisciplinary discussions. The aforementioned canvases, with modifications made based on the participants’ feedback, are presented in Section 4.

In the remainder of this article, we will discuss the limitations and advantages of existing design frameworks and describe in detail the implementation process of our framework in the two case

Fig. 1. The proposed framework for the design of social robots, with the three phases of the design process: (1) definition of the problem space, (2) creation of the design guidelines, and (3) integration of them in the solution space.

studies. In addition, we will analyze the feedback of people who participated in the two case studies and discuss the implications that the use of this framework can have for future research and for the successful development of social robots.

2 PARTICIPATORY DESIGN AND DESIGN FRAMEWORKS

We present a design framework, which can be used as a tool during the **Participatory Design (PD)** of a social robot. PD is a practice of **User-Centered Design (UCD)**, where multiple stakeholders contribute to the design process as active co-designers [12, 67, 88, 105]. While the definition of PD sometimes focuses solely on involving users alongside designers [88], we broaden the definition of PD as involving also other stakeholders, such as domain experts (e.g., in the case of a robot designed for use in autism therapy, a psychologist specialized in autism would be considered a domain expert) [12, 105].

The aim of PD is for designers to incorporate the realities of the stakeholders' situation into the design process, while for stakeholders the aim is to be able to articulate their needs and goals, and to find appropriate technological means to achieve them [88, 98]. PD has the advantages of examining users' needs, challenging assumptions, encouraging reciprocal learning, and creating new ideas [36, 67]. By treating stakeholders as designers rather than consumers, the design team can boost its ability to frame and solve problems, and create value [36, 61]. Major issues of concern in PD include taking into account the domain specialists' expertise, beneficial and sustainable innovation, taking seriously multiple viewpoints and the facts and resources they provide, taking into account the context of the product, the authentic experience of the experts regarding the problem being solved, hands-on methods to solve real world problems, and reflective practice of design [98]. The practice of PD in the field of HRI will be examined in Section 2.1.

A design framework—in the context of this article—is defined as a structure that underlies and supports the process of designing. In general, frameworks are models used to describe complex

phenomena, and the relationships between the parts that construct those phenomena. In the field of HRI, frameworks have been previously used to describe, for example, the various permutations of “Oz” and “Wizard” states in HRI [90], and the degree of sociality between robot and user [50]. For a more detailed discussion about design frameworks in HRI, see Section 2.2.

Design frameworks are separate from tools for PD, as we do not require design frameworks to be tangible in our definition: They can also exist as mental models. For this reason, we regard our contribution as both a design framework (descriptive model of a certain design process), as well as a tool for PD (tangible tool intended explicitly for multidisciplinary design teams), and we conduct a review of both tangible PD tools and design frameworks in the field of HRI.

2.1 Practice of Participatory Design in Human–Robot Interaction

The use of PD has been argued for in the field of Science and Technology Studies, due to the inseparability of the social context and the technological material, that is, the practice and the technology [92]. All technological systems exist in the context of a society—they are in effect always socio-technical systems. For this reason, both the experts within the system and the users of the system have valuable information with which they can contribute to the design of those technological artifacts, and will be the ones who are affected most by changes in the system [68, 92]. PD emphasizes power dynamics and hierarchies between designers and users—and aims to share decision-making power with stakeholders, encourage critical discussions about the meanings tied to the technology being developed, as well as allow stakeholders to learn design and technology know-how from designers [57, 88].

In HRI, PD has been used to design robots for different types of target populations (e.g., older adults with depression [57] and dementia [64]) and for different interaction scenarios (e.g., a service robot for blind people [6], and robot-assisted therapy for children with autism [5, 48]). Within this multitude of different contexts and targets, the approaches used in PD for HRI have been equally varied.

Broadly, all design processes loosely follow the same steps. In particular, the authors start by (1) gathering some initial information (interviews, surveys, etc.). After that, they (2) share that initial information among participants as a base for the design; and finally they use the information gathered to (3) design a robot that is adapted to a particular set of users or scenario.

For example, Lee and colleagues described an instance of collaboration among a group of target users (older adults) and experts (therapists) to design a robot intended to interact with older adults with depression [57]. The researchers initiated collaboration among experts and co-designers by performing interviews to gather background information, and then by motivating discussion by introducing the basics of social robots, sketching robotic solutions, thinking of requirements for a robot, physically prototyping the robot, and organizing separate design sessions with domain experts.

In addition, PD has also been used to design robots to support people with dementia. Moharana et al. described informal caregivers as one of the most important stakeholders in dementia care, and involved them in the co-design process [64]. The process employed co-design methods such as performing interviews to collect background information, storyboarding, physical prototyping, group ideation, and scenario cards.

In the work of Azenkot et al., a robot to guide blind people was designed with the PD approach by three designers and five non-designers, with all but one designer having a vision disability [6]. For this design, an existing robot was used, and new interactions were designed both during the course of a group storyboard session, as well as individual design sessions with a designer and a prospective user testing the robot.

PD was also used to design HRI interventions for children with autism, together with autism therapy professionals, parents of children with autism, and adults with autism [48]. This process did not strictly design the robot itself, rather the interaction scenarios to be completed with it. Another result of the co-creation process was a framework for designing new interventions with robots for children with autism. Similarly, PD was used by Axelsson et al. to design a robotic tutor of sign language for children with autism, together with autism therapy providers and roboticists. The robot was tested in a user study by children with autism, where feedback on the design was provided by the children and their companions [5].

2.1.1 Tangible Participatory Design Tools in Human–Robot Interaction. Tangible tools for the PD of social robots and HRI are limited. The tools that are available, which are introduced here, have been developed for specific application domains within HRI, and are not applicable flexibly to different contexts.

In the application of a robot helping older adults with depression, the PD team involved in building the application utilized cards that pictured the older adults' daily activities [64]. The cards were used to facilitate conversations about the challenges the older adults encountered related to their daily activities. After this, the design team discussed how a robot could be used to address these problems, and selected cards to describe the future robot's interaction modalities. The researchers present the cards as a potential tool for building similar future robots [64].

Huijnen et al. developed an intervention design template, to aid the design of HRI therapy interventions for children with autism [48]. The template was developed by eliciting requirements from professionals and adults with autism. The template, which can be printed onto paper and thus worked on together by the design team, helps the designers define the requirements of the robot, its end-user, its environment and practical implementation. This systematizes the description of robot interventions in autism therapy [48]. This tool will be described in more detail in Section 3.1, as it was especially influential to the development of our framework.

A social robot design toolkit was created with the aim to teach programming concepts to preschool children, by embedding programming knowledge into an interpersonal interaction with a social robot. The kit does not strictly aim to design a social robot or HRI, but is notable in this context due to its tangibility and incorporation of children into the design process of a robot's behaviour. It uses tangible materials such as vinyl stickers, which are familiar to preschoolers [44].

All of the aforementioned studies have in common the involvement of stakeholders such as domain experts or users as co-designers. All of the studies also mention workshops or group brainstorm sessions as methods for co-creation [5, 6, 48, 57, 64]. A few also mention using domain-specific frameworks or tools to facilitate the group sessions [48, 64]. However, the studies make no mention of tools applicable to different contexts for the purpose of PD of HRI being available. This lack of easy-to-use tools to use in the process of designing social robots hinders the design process of multidisciplinary teams. This is due to participants in the process of design might struggle to materialize the different design and data collection steps, and to reconcile knowledge and contributions coming from different academic backgrounds. It is important to have PD tools for HRI available so that researchers can use them for co-design with experts across domains, and prospective users. The framework we propose addresses this issue. The main strengths of our proposal come from the malleability of our framework, in other words its wide applicability to multiple HRI design problems—and its potential to introduce new ideas and considerations to different domains within HRI, where they are not typically considered. It also takes into account safety and ethical considerations, which has been encouraged by the robotics community [104] but not explicitly incorporated in other frameworks. These properties are discussed in more detail in the next section.

2.2 Design Frameworks in HRI

New media—such as frameworks—can be used to support social creativity that occurs when co-designers have a “symmetry of ignorance”—i.e., they have expertise in different fields [36]. In the design of technology, previous design frameworks have been used to support end users’ viewpoints and carrying out collaborative design. For example, *design thinking* [20, 59] is a methodology prevalent in Silicon Valley. It is characterized by pursuing an understanding of and empathizing with the user, redefining problems, and creating ideas and prototypes through iteration. *Meta-design* [38], which has been used to support domain experts to act as designers in software development [37], is another example of such a framework. It defines activities, processes and objectives for creating new media that allow users to act as designers. While these methodology frameworks can also be applied in the design of social robots, they introduce no physical tools, and do not introduce the knowledge on social robots that the framework proposed by the authors does. This is a limitation that stands in the way of their implementation in real-life scenarios as technology developers often lack the expertise necessary to materialize these frameworks.

Design frameworks for HRI have previously taken several forms, which range from conceptual to computer software supporting the development of the robot’s software. For example, frameworks can take the form of tangible software programs. A framework describing different interaction modes as skills for the social robot Maggie was developed by and for researchers working on the robot [45]. The skills employed tactile, visual, remote voice and sound modes. In another project, roboticists introduced a component-based design framework for robot software architecture, with the aim of facilitating developers of the software [107]. A graphic interface for developing interaction sequences for social robotics was implemented to allow programmers and interactions designers to work together [42]. The authors present the interface as a design framework that increases the effectiveness of programmer-designer teams developing social robot applications.

Frameworks can also be entirely conceptual, where a structure or hierarchy of design considerations is presented to help the design process. A conceptual PD framework for developing studies for robots used in education is focused on three components: Time, Space, and Structure [11]. The framework was evaluated in a case study where the therapeutic robot seal Paro was used as a platform for applications developed for children with autism at school. Design teams consisted of teachers of the children, and the children themselves. In another study, roboticists collected a list of values from experienced humanoid roboticists via interviews, and advocated for their consideration during HRI design, as an example of the Value Sensitive Design method [24].

Frameworks can also be have a physical form to express the design structure and hierarchy, but not necessarily support the creation of software. The framework presented in this article falls into this category. Previous examples of such expressions of frameworks are for example a design framework for mobile robot systems that is expressed in the Unified Modelling Language, primarily using class diagrams [43]. The framework aims to define the functionality of many common robot system components. A design framework by Bartneck and Forlizzi describes the design dimensions of a social robot on a general level and is applicable to multiple contexts, however the authors point out that a framework with more detail should be developed [8]. This framework will also be described in more detail in Section 3.1, as it was highly influential to the development of our framework.

3 FRAMEWORK DEVELOPMENT

In this section, we present which previous design frameworks and tools influenced us in the creation of the framework, how we iteratively designed the proposed framework, and how we examined it in co-design sessions.

3.1 Influential Design Frameworks and Tools

In the creation of the design framework we were mainly influenced by two previously proposed design frameworks for social robots [8, 48]. Our aim was to incorporate both the broad design context introduced by Bartneck and Forlizzi [8], as well as the detailed description of interaction scenarios (specifically designed for use as a tangible tool by multidisciplinary teams) described by Huijnen et al. [48]. Both frameworks have also been demonstrated to be usable in a real world context: Bartneck and Forlizzi's design framework has been applied to analyze existing robots, and Huijnen et al. have created autism therapy interventions with the help of their tangible framework tool. This makes them especially influential, as the framework presented here is also designed to be usable in a real world context, rather than just exist as a theoretical tool.

The framework proposed by Bartneck and Forlizzi [8] describes social robots through five dimensions: *form*, *modality*, *social norms*, *autonomy*, and *interactivity*. After describing these dimensions, the article defines guidelines for social robots' design. The authors call for the future development of a more detailed framework and more specific guidelines, as robots are designed and built to address specific application areas [8].

The framework proposed by Huijnen et al. [48] is described as an intervention template, specifically developed to design interventions for children with autism, with the robot Kaspar or other robots [31]. The framework was developed on the basis of insights gained from children's parents and experts of autism treatment, as well as from the authors' previous work. Their template is used as a tool to aid co-creation sessions with the goal to design new interventions for robots used in autism therapy. The design dimensions used are *robot*, *end-user*, *environment*, and *practical implementation* with the creation of a *new intervention* as a result of the design process.

Our proposed design framework improves upon and broadens the scope of use of the aforementioned frameworks. This framework provides the level of detail suggested by Bartneck and Forlizzi [8]. The framework by Huijnen et al. [48] specifically designs interventions but not robots to perform those interventions with. Our framework combines the broad scope and flexible context of Bartneck's framework, as well as the level of detail and function as a tool of facilitation during a design process of Huijnen's framework.

3.2 Research Through Design and Iterative Design Approaches

We developed a framework for the PD of social robots with a **Research Through Design (RtD)** approach, a design paradigm used in **Human-Computer Interaction (HCI)** [108]. RtD highlights the role of the designed artifact as the chief element in the process of generating and communicating knowledge—in this case, that artifact is the introduced tool for PD. The approach has the goal of creating the “right” thing, to transform the world from its current state into its preferred state. Additionally, we used the iterative design approach [71], where each version of the tool was evaluated for usefulness after each improvement. We argue that the criteria for useful RtD—detailed and reproducible process, significant invention combining information in a novel manner, relevance to the field, and extensibility to future design problems—are fulfilled by the PD tool [108].

Our framework was first envisioned during the design process of a social robot to be used as an assistive sign language tutor for children with autism spectrum disorder [5]. Three roboticists and three domain experts of autism therapy—a speech therapist, a neuropsychologist, and a health care quality manager—took part in the PD process. The initial design dimensions and variables detailed in the framework arose from previous literature on social robots, as well as the initial co-design session creating the robotic tutor of sign language for children with autism.

We then placed these dimensions and variables onto canvases to make them into a boundary object (discussed in more detail in Section 3.3) that functioned as a tool for PD. We then

iteratively improved upon the tool, to make it more understandable, broad and exhaustive in scope, and suitable for co-design by a multidisciplinary team. We employed a communal approach to the refinement of the tool, since the tool is designed for the use of multidisciplinary teams, and thus feedback from multiple people across different fields was considered important. The improvements were based on data gathered from 7 workshop sessions, with 97 unique participants across multiple disciplines using the tool and giving feedback on its use. The resulting tool is presented here.

The first version of the tool was tested by a designer, developer and roboticist, to create a design concept for an office robot. It consisted of three canvases, one for each phase of the framework (problem space, design guidelines, and solution space). The feedback called for more detailed canvases, separated into clearer modules. After this, the framework was expanded into seven named canvases: (01) Problem space, (02) Ethical considerations, (03) Design guidelines, (05) Environment, (06) Form, (07) Interaction, and (08) Behaviour. The canvases (05)–(07) represent the solution space, now divided into four clear dimensions.

These canvases were evaluated with two participants (a developer consultant and an analyst consultant). After this session, canvas (04)—**Minimum Viable Product (MVP)**—was created due to the request of the participants. The participants regarded canvases (05)–(08) as important, but thought that a summarized version of these canvases would be useful for when they had limited time. The canvas (04) Minimum Viable Product is named after a term common in product development, where an MVP version of a product is developed initially with only “bare bones” features. This is done because an MVP takes less time to develop, and can be quickly deployed to test its feasibility [65].

The canvases (01)–(04) were next evaluated by a group of 40 participants across multiple disciplines, attending a workshop where design concepts of social robots were created. After this, the same canvases were evaluated by 30 design students. These canvases were improved based on feedback after each session. Improvements on the canvas (04) MVP were also made to the more detailed solution space canvases (05)–(08).

After these iterations, all canvases (01)–(08) were systematically piloted in two projects, where the aim was to build a robot based on the created designs. The first robot was designed to play games online with teenagers for a Finnish public media company, with roboticists and media specialists from the media company designing together. The second was a robot to guide library customers to books and book categories in Helsinki central library Oodi, designed together by roboticists, librarians, and library users. After both projects, feedback on the current canvases’ advantages and limitations was collected by giving feedback forms to participants, with questions about the canvases’ efficacy. The feedback from these two projects is presented in this article, and is used to illustrate the advantages and limitations of the framework.

After the collection of this feedback, modifications were made. The canvases were then reviewed by 5 roboticists, and final modifications were made. Two additional canvases (09) Service ecosystem and (10) Experience flow were created as tools to supplement the design process, due to participants requesting tools to think about the robot’s operation context even more broadly. This version of the canvases (01)–(10) is presented here. Feedback from the two projects (designing a gaming robot and a library robot), which resulted in the building of real robots is introduced here—consisting of 15 unique respondents including roboticists, domain experts, and users.

3.3 Using the Canvases as Boundary Objects

Frameworks for design—when made tangible—can act as externalizations or *boundary objects*, which help co-designers articulate their thoughts in a shared language, and to incrementally create a shared understanding of the design problem and an appropriate solution [36, 70]. Boundary objects are characterized by their ability to serve as bridges between intersecting social and cultural

worlds [70]. The canvases can also be used to support iterative design, a method often used in HCI design [71], and track the history of proposed changes.

Canvases have been previously used as boundary objects for PD. The *business model canvas* is a design tool that aims to structure sustainability issues in business model innovation [49]. The *Lean Service Creation canvases* are a set of open source [40] canvases, which aim to aid the collaboration of teams where members have different ways of working and skills, by bringing a common language to the designers [74].

The presented framework can be utilized as a boundary object when printed out as *canvases*, either as A1 or A3 paper sizes—being big enough for all participants to read and work on simultaneously, by writing or applying post-its to them. 10 canvases are introduced, each with its own role within the PD process, addressing a different design issue.

4 FRAMEWORK FOR PARTICIPATORY DESIGN OF SOCIAL ROBOTS

This section introduces the entire framework, section by section. The framework is divided into three phases, as shown in Figure 1: (1) the problem space, (2) design guidelines, and (3) solution space. Figure 2 illustrates what the team can achieve with each canvas, and their recommended sequence of use.

4.1 The Problem Space

The problem space of our framework describes the problem to be solved in detail. It encompasses two canvases: (01) Problem Space and (02) Ethical considerations.

4.1.1 (01) Problem Space. The problem space, pictured in Figure 3, examines the problem to be solved, both from the *user's* and *robot's* perspectives.

Users are defined through their group(s), characteristics, needs, and goals. Focusing on what characterizes the user, and what they are trying to achieve, is a central method in UCD [1], which aims to turn needs into design requirements. Both primary and secondary users are examined (e.g., primary being students using the robot in a classroom, and secondary being teachers helping students use the robot). Goals are examined from short and long-term perspectives, as users' needs can change during long-term interaction with the robot [64]. The explicit definition of a goal helps examine the eventual solution's success. Using the challenges the users face as a starting point for PD has been previously used in HRI [57].

The robot is defined through its task(s) and advantages. The task of the robot responds to the user's goal. It is examined in terms of short and long term.

A robot can have several advantages when compared to using a human or another technical device for the solution. Examining what explicitly are the robots' advantages can help avoid building applications when there is no specific benefit from using a robot. The list of advantages in the framework is not comprehensive, but aims to give the design team perspectives of a robot's potential advantages: social skills [16], the user's emotional response [106], personalization of interaction [60], precise tasks [96], active data collection with sensors [19], mobility [81], environmental manipulation [72], and the ability to connect to other technical systems [95].

4.1.2 (02) Ethical Considerations. The ethical considerations canvas is pictured in Figure 4. The target of robot ethics is the human ethics of the robot's designers, manufacturers, and users—rather than building artificial notions of ethics for the robot. Ethical considerations emerge from the user and the robot, guiding the interactions between the two [104]. Engineers are “materializing morality” when they design technologies, as the technologies shape human actions, actions that could be considered moral or amoral [103].

Fig. 2. Instructions of when to use each canvas. Two separate design paths are depicted. Path 1 is selected when the teams have limited time and want to make a quick first draft of the robot's design. Path 2 is selected when the team wants to create an in-depth design of the robot. Canvases (09) and (10) are optional. The phases of the design process are depicted here, similarly to Figure 1.

By explicitly including ethics in the problem space, designers can attempt to protect users from potential negative effects of the robot. Value Sensitive Design is the proactive integration of ethics during the technology's design, which accounts for human values in a principled and comprehensive manner throughout the design process [39]. While no list of values is ever exhaustive, this

Canvas # 01

PROBLEM SPACE

SOCIAL ROBOT CO-DESIGN CANVASES

What problem are you solving?

USER

Group(s) <small>Who are the users? Are there supporting users? For example: students and teachers.</small>	Characteristics <small>What are the users like?</small>	Needs <small>What do the users need?</small>
primary users		
secondary users		

Goal(s) <small>What do the primary and secondary users want to accomplish?</small>	
primary users	
secondary users	
SHORT-TERM	LONG-TERM

ROBOT

Task(s) <small>What task does the robot perform?</small>	
SHORT-TERM	
LONG-TERM	

Advantages <small>What advantages does using a robot bring (compared to a computer or human)?</small>			
Social skills	User's emotional response	Personalization	Precise tasks
Data collection with sensors	Mobility	Environment manipulation	Connection to systems

Social Robot Co-Design Canvases free version by Minja Axelsson is licensed under a Creative Commons
 Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) license. Sponsored by Futureice.
 The Social Robot Co-Design Canvases can be found at <https://osf.io/jg2t8/>
 Cite as: Minja Axelsson. 2020. Social Robot Co-Design Canvases. <https://doi.org/10.17605/OSF.IO/JG2T8>

Fig. 3. The problem space canvas.

design framework provides a list of often raised ethical considerations relevant to social robots: physical safety, data security, transparency, equality across users, emotional consideration, and behaviour enforcement.

Physical safety—The physical safety of the user needs to be ensured while using the robot [41]. Machinery has the potential to crush or pinch the user if operated recklessly. As for Physical HRI, appropriate touching protocols are needed to provide psychological safety.

Data security—The safety and privacy of data collected by social robots is an ethical concern [22, 26]. Robots are in a unique data collection position when compared to personal computers, because they can act as social agents and elicit emotional responses, which may lead to revelation of larger quantities of personal data by the user, as well as data of a more personal nature [27]. If robots are used for tasks where the user is especially vulnerable, for example bathing and sanitation in elderly care, then it is recommended that data collection be switched off for these tasks [76]. Tradeoffs between functionality and privacy need to be carefully considered [27].

Transparency—The transparency of an autonomous system can be defined as the sharing of accurate perceptions of its abilities, intentions and constraints with the end users [62]. Bartneck and Forlizzi [8] define two guidelines in their framework that aim for two facets of transparency:

Fig. 4. The ethical considerations canvas.

“The form of the social robot should match its abilities” and “Transparency of capabilities through physical appearance.” They also mention how robots should communicate their internal states accurately. A human should be informed of what purpose the robot was built for, and what is the intent of the system [62, 76, 97]. If a person overestimates the capabilities of a robot due to lack of transparency, then this can lead to misuse or over-reliance on the robot. A reliable perception of the robot’s capabilities will help the user form accurate perceptions of the robot’s ability, intent and situational constraints [62].

Transparency can become a tradeoff with utility [97]. For example, in a Wizard-of-Oz experiment setting, immediate real-time transparency can be detrimental. However, not disclosing Wizard-of-Oz methodologies may lead to inappropriate expectations of robots [76]. Different users may also require different amounts of information (e.g., children vs. adult users), and research into domain-specific appropriate levels of transparency may be needed [97].

Equality across users—Stereotyped form and behavior of robots can negatively affect the attitudes that humans have toward other humans. Harmful racial and gender biases may be going unnoticed among developers and users of robotic technology [27]. For example, the construction of robot identity through naming has been shown to reveal gender biases, wherein male robot

names expressed mastery by referencing for example to Greek gods, while female names tended to be infantilizing or sexualizing [51].

There is also a lack of diversity in robot morphology and behavior. The vast majority of humanoid robots have Asian or Caucasian features, with tendency to conform to masculine stereotypes [76] or overly feminized robots [78]. Inequality can also appear in machine learning software used to control the robot's behavior if the datasets used for training are biased [23]. Designers should avoid racist, sexist, and otherwise harmful or unequal qualities both in the robot's form and its behaviour [76].

Emotional consideration—Humans have a subconscious tendency to treat robots as if they were alive, even if they are clearly not [26, 35, 75]. An anthropomorphic appearance can cause people to consider robots more lifelike [82], and giving it a backstory, name, and character can cause people to be more empathic [28]. Social robots' design should thus accommodate human emotions, by considering whether anthropomorphism should be encouraged or discouraged [27, 35, 76]. The domains of elderly care, childcare, health care, and education are of special importance due to the vulnerability of the users in these domains [26, 76]. Designers should also evaluate whether the development of a relationship with the robot before, after, and over several interactions is beneficial, or ethical [32].

Behaviour enforcement—There is evidence that negative behavior toward robots has the potential to traumatize, or desensitize users to negative behavior toward humans [26]. Both children and adults have shown potential to treat social robots in an abusive manner when unsupervised [21, 52]. Designers should consider how to minimize negative behavior toward the robot, for example with negative user feedback (e.g., the robot discouraging the user verbally or with body language from unacceptable behavior).

4.2 Design Guidelines

Design guidelines or design principles have often been previously used in HRI to create robotic applications on different abstraction levels: guidelines for specific robotic applications (e.g., design guidelines for a peripheral robotic conversation companion [47] or robots to care for elderly people with depression [64]), guidelines for robot behaviours (e.g., proxemic behaviours for the robot PR2 [94]), and general HRI guidelines (e.g., Reference [8]).

Canvas (03), pictured in Figure 5, formalizes the creation of design guidelines. It encourages the design team to reflect on the advantages considered on the canvas (01) Problem space, and the ethics on the canvas (02) Ethical considerations. Next, the design team can consider what guidelines could be suitable for the dimensions of the robot's solution space: (05) Environment, (06) Form, (07) Interaction, and (08) Behaviour. The guidelines provide a method of abstracting the information considered in the problem space onto the solution space of the robot.

4.3 The Solution Space

The solution space considers four dimensions of a social robot: environment, form, interaction, and behaviour. Each dimension is examined through different qualities.

4.3.1 (05) Environment. The environment canvas, pictured in Figure 6, examines factors surrounding the robot's operation. While not all factors can be influenced by the robot's designers, becoming aware of them during the design process can help the designers better adapt the robot to its environment—and plan any changes that are possible.

Where and *when* describe the setting where the HRI happens.

Users and *secondary users* relate back to those defined 4.1.1. Here, they are brought back to mind as components of the solution space. Secondary users can act as facilitators of the interaction

Fig. 5. The design guidelines canvas.

between the primary user and the robot, such as in Reference [41]. As mentioned in the ethical guidelines, the involvement of another human in the interaction process should be considered especially important in domains that involve emotionally vulnerable users, such as health care, child care, elderly care and education.

Simultaneous users describes the number of users simultaneously interacting with the robot [55]. On the canvas, this is defined from one, to few, to several. The canvas indicates a tradeoff of more simultaneous users requiring a more sophisticated robot, to help the design team balance between number of users and development time.

Next, technical components of the environment are introduced. *Data collection* from the environment is considered [19]. A tradeoff between data collection and a required investment in data security is introduced.

External sensors and actuators, which are placed in the robot's environment as part of its system such as in Reference [10], are considered. Finally, *connection to other systems* such as databases is considered [95].

4.3.2 (06) Form. The form of the robot describes the robot's outwardly perceptible qualities. The canvas is pictured in Figure 7.

Canvas # 05

ENVIRONMENT

SOCIAL ROBOT CO-DESIGN CANVASES

What is the robot's context of operation?
You can use the "Ecosystem" canvas to dive deeper into this topic.

<p style="text-align: center;">Where</p> <p style="text-align: center;">What place? Does it change?</p>	<p style="text-align: center;">User(s)</p> <p style="text-align: center;">Who is using the robot?</p>
<p style="text-align: center;">When</p> <p style="text-align: center;">What time of day? Does it change?</p>	<p style="text-align: center;">Secondary user(s)</p> <p style="text-align: center;">Are there secondary users? E.g. teachers that help students use a robot.</p>
<p style="text-align: center;">Data collection</p> <p style="text-align: center;">Does the robot collect data from its environment? How is it stored?</p> <p style="text-align: center;">⚖️ TRADE-OFF: More data collection requires more attention to data security.</p>	<p style="text-align: center;">Simultaneous users</p> <p style="text-align: center;">How many users should be able to use the robot simultaneously?</p> <p style="text-align: center;">⚖️ TRADE-OFF: More simultaneous users requires a more sophisticated robot.</p> <div style="text-align: center;"> </div>
<p style="text-align: center;">External sensors and actuators</p> <p style="text-align: center;">Does the robot use external sensors? Does it have external actuators, such as lights or limbs?</p>	<p style="text-align: center;">Connection to systems</p> <p style="text-align: center;">Is the robot connected to external systems, such as software, databases, or other robots? How does it use these systems?</p>

Social Robot Co-Design Canvases free version by Minja Axelsson is licensed under a Creative Commons
 Attributions-ShareAlike 4.0 International (CC BY-SA 4.0) license. Sponsored by Futurice.
 The Social Robot Co-Design Canvases can be found at <https://osf.io/jg2t8/>
 Cite as: Minja Axelsson. 2020. Social Robot Co-Design Canvases. <https://doi.org/10.17865/OSF.IO/JG2T8>

Fig. 6. The environment canvas.

A well-known measure of a robot's form is evaluating the robot according to its "familiarity" to life. This is defined by the degree to which the robot resembles humanlike appearance and movements. It is argued that a robot should not be too humanlike, or it will fall into the "uncanny valley" of negative familiarity, appearing like an animated corpse. Instead, a safe familiarity can be produced by a design approaching human likeness, but not too close to it [66]. This theory can be further extended to zoomorphic robots: They should not be made too lifelike, to avoid appearing zombielike. All dimensions of the robot's form contribute toward its lifelikeness, with appearance and movement being most important.

The design team is asked to draw a picture of the robot to get their imagination flowing and to create a common reference point. This method has been previously used in PD of HRI [57].

The *appearance* of a robot is highly variable, which is why a two-dimensional map is used to describe it. Informed by previous work [8, 22], we define the two dimensions going from machine-like to lifelike and from anthropomorphic to abstract. The appearance of a robot defines the user's initial response to it. If a robot appears sophisticated, then the user will assume a similar level of sophistication in its skills [8, 30]. A humanoid robot is perceived as humanlike and elicits strong expectations about the robot's social and cognitive competencies [30], while

Canvas # 06

FORM

SOCIAL ROBOT CO-DESIGN CANVASES

What are the robot's outward qualities? If an existing robot is used, are its qualities modified?

Draw a picture

What does the robot look like?
Is it attached to something?
Does it move around?
Can its appearance be modified?

Appearance

Is the robot more machine or lifelike? Is it human-shaped, animal-shaped, or abstract?

TRADE-OFF:
Robots that appear more human and lifelike are expected to be more sophisticated in features.

Size

How big is the robot?

Character of movement

What is the robot's movement like?

Mobility

Does the robot move across space? Does it move in place?

Voice & sounds

Does the voice have a gender or an age? What are pitch, speed and prosody like? Is the voice always the same?
Does the robot make sounds: music, "beep"s, animal noises?
When are these sounds heard?

Visual cues

Does the robot have expressions, lights, a screen or other visual elements?

Touch & smell sensations

Is the robot soft or rough, warm or cold?
How does the robot smell?
Touch and smell are especially important in close interactions.

Social Robot Co-Design Canvases free version by Minja Axelsson is licensed under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) license. Sponsored by Futurice.
The Social Robot Co-Design Canvases can be found at <https://osf.io/jg2t8/>
Cite as: Minja Axelsson. 2020. Social Robot Co-Design Canvases. <https://doi.org/10.17605/OSF.IO/JG2T8>

Fig. 7. The form canvas.

an abstract robot is perceived as more of an animal, with a lower level of functioning. Zoomorphic robots fall in the middle of the scale, which ranges from anthropomorphic to not humanlike at all.

The *size* of the robot is rated on a scale from smaller than (the average adult) human, to human-size, to bigger than human. The size of a robot can affect its usefulness within a certain task [86].

The *character of movement* of a robot can be rated from machinelike, to hybrid, to lifelike. Lifelike, smooth movements suggest a more sophisticated robot in terms of skills, while a more mechanical movement quality suggests a more simple robot [34, 80].

The robot's *voice and the sounds* contribute to its lifelikeness [15]. A voice can be robotic, or humanlike. Voice can be perceived to have a gender or an age, which depends on speed, pitch and other qualities. Voices, both robotic and humanlike, can have tonality, being either monotone or having vibrant prosody. The noises a robot can make are almost infinitely variable: machinelike "beeps," music, humanlike exclamations or animal calls are common examples of sounds a robot can make. The sounds that a robot's machinery may make when it moves are also part of its soundscape. A robot's soundscape can also include sounds that it makes when it is touched or, if it is mobile, sounds that it makes when it moves in space.

Canvas # 07

INTERACTION

SOCIAL ROBOT CO-DESIGN CANVASES

How does the robot interact with the user?
You can use the "Experience Flow" canvas to dive deeper into this topic.

Interaction modalities

What modalities are inputs to the robot? What modalities does the robot output?

INPUT	OUTPUT
☐ voice	☐ movement
☐ sounds	☐ touch
☐ gestures	☐ screens
☐ smell	☐ lights
☐ facial expressions	☐ other _____
☐ other _____	☐ voice
	☐ sounds
	☐ smell
	☐ gestures
	☐ facial expressions
	☐ other _____

Interaction flow

Describe the most important interaction of the robot.
Note: only fill the bottom row if your robot is teleoperated.

	BEFORE	DURING	AFTER
USER			
ROBOT			
ROBOT OPERATOR (optional)			

Situation flow

How defined is the situation where the interaction takes place?
Does the user always enter and exit at the same point?

predefined
flexible
freestyle

Leadership

Who initiates the interaction? Who determines what happens next?

robot-led
mutual / alternate
user-led

Goal

What is the user's goal in the interaction? What describes the interaction?

task completion / informative
both
explorative

Robot's name

Does the robot have a name which is used during interaction?

⚖️ **TRADE-OFF:**
A robot with a name, creates more emotional bond.

Social Robot Co-Design Canvases free version by Minja Axelsson is licensed under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) license. Sponsored by Futureice. The Social Robot Co-Design Canvases can be found at <https://osf.io/jg2t8/>.
Cite as: Minja Axelsson. 2020. Social Robot Co-Design Canvases. <https://doi.org/10.17605/OSF.IO/JG2T8>

Fig. 8. The interaction canvas.

Mobility is described within the form of the robot, because the way a robot moves affects its embodiment. The degree of mobility of a robot can affect user perception toward it [93].

Visual cues, such as lights [7] or a screen [102], can be added onto the robot's form.

Touch sensations are especially important for Physical HRI. Soft and warm tactile sensations suggest comfort and familiarity, while hard and cold tactile sensations suggest distance and industrial qualities [2]. *Smell sensations* should be considered [8] and are also especially important if the user is in close contact with the robot. Olfactory sensations can induce sensations of pleasure or disgust, or they can be entirely neutral.

The selection of an existing robot for a solution usually brings constraints that impact our definition of form. However, modifications can be made to alter things such as voice or even appearance by having the robot wear a costume [79].

4.3.3 (07) Interaction. The interaction dimension, pictured in Figure 8, defines the manner in which a user interacts with a robot.

Interaction modalities define the different ways a user can interact and communicate with the robot. Input and output of the robot are examined separately. Modalities of interaction can vary greatly [10, 73, 101]. A checklist of the most common modalities encountered in interaction with

robots are listed in this framework: voice, sounds, gestures, movement, touch, smell, facial expressions, screens, and lights. Additionally, an empty space for other modalities is provided.

Interaction flow allows the design team to formulate the basics of the planned HRI: what the user and robot are doing before, during, and after the interaction. Additionally, the team can fill in what a robot operator is doing if the robot is being operated by a human.

The *situation flow* describes how predefined the interaction is, ranging from predefined [13], to somewhat defined [69], to freestyle [55]. Situation flow describes to what extent the events leading up to and after the interaction are defined and is thus different from the interaction itself.

The *leadership* of the interaction describes who drives the interaction forward and takes initiative [46]. Leadership can range from robot-led to mutual to user-led. Leadership defines who is in control of the interaction, e.g., who is initiating it or dictating its pace through turn-taking [85].

The *goals of the interaction* relate back to the goal defined in the problem space. The interaction can have one or multiple goals. We describe the goals based on two extremes: The goal is the completion of task (some condition is satisfied or some procedure is completed) [10], or the goal is the interaction itself [55, 56]. An exploratory interaction can for example generate knowledge, pleasure, or a novel experience for the user.

Finally, if the *robot has a name* that it answers to during the interaction, then the design team can fill it in. However, giving a robot a name and a “character” can cause people to be more empathic toward it [28], which is also an ethical consideration. This tradeoff is introduced on the canvas.

4.3.4 (08) Behaviour. The behavior dimension, pictured in Figure 9, defines how the robot acts. Behavioral design is important, as the robot’s behavior is one of the primary determinants of the user’s attitude toward it, and contributes to its lifelike impression [30].

The *role* of a robot can be highly variable and has been explored in several studies, e.g., References [3, 46, 48]. Huijnen et al. [48] defined seven robot roles for autism therapy: “provoker,” “reinforcer,” “trainer,” “mediator,” “prompter,” “diagnoser,” and “buddy.” However, for the design of social robots flexible to multiple contexts, these roles are not sufficient. The design team is welcome to write down the role they feel describes the robot best.

Motivation of behavior describes whether the robot behaves in response to external stimuli from its environment, purely according to internal models, or a mix of the two. An example of motivation via internal models is the motivation system of the robot Kismet, in which behaviors are influenced according to internal models of Kismet’s drives and emotions [18].

The design team is invited to consider the robot’s *personality*: whether it has specific characteristics, needs, or emotional states. A robot’s personality [77], e.g., its level of extroversion [58], can affect how users respond to it.

Mode of operation defines to what degree a robot behaves without human input. Autonomy of the robot’s behavior can range from fully human-controlled, partially human-controlled, to fully autonomous [9]. For example, a robot that is fully human-controlled has no autonomy while a robot with partial autonomy may occasionally need human intervention to act in its environment.

Social skills describe how well the robot follows social conventions, such as greeting a new person when they enter a room, maintaining appropriate personal distance to a human, or understanding when it is its turn to speak. Designing social behavior is difficult, as a robot may be perceived to have intentional behavior, due to its lifelike appearance and behavior [8]. Social robots should adhere to generally accepted social norms to create the impression of possessing some form of social intelligence [8, 30]. The tradeoff of extensive social skills requiring a more sophisticated robot is introduced. The team then defines in detail what *social behaviours* the robot exhibits.

Contextual adaptation of behavior refers to how much the robot can independently determine its behavior according to the context. Context can include environmental variables, such as weather

Fig. 9. The behaviour canvas.

and time, as well as variables of the user, such as mood. The levels of contextual adaptation of a robot's behavior range from none [13, 55] to partial [63, 87] to extensive. Robots that exhibit no contextual adaptation always follow the same course of action in a given situation. Their behavior is "hard-coded." Behaviour that can be classified as extensive contextual adaptation would involve the robot taking input from past situations that it has learned from, as well as the current context and determining the best course of action according to these. In this case, the robot would have a "memory." An example of partial contextual adaptation could be a robot behaving in a structured way (e.g., following a script) but with randomized options or context specific choices in its structured behavior. The design team then defines what *contextual behaviour* the robot displays in detail.

Personalization means the robot adapting its behaviour toward a specific user [60]. The tradeoff personalization requiring more personal data to be stored is introduced, which is also an ethical consideration (as introduced in Section 4.1.2).

4.4 Additional Canvases

These three additional canvases were made due to participant feedback requesting additional tools to examine particular aspects of the solution space.

Canvas # 04

ROBOT DESIGN MVP

SOCIAL ROBOT CO-DESIGN CANVASES

It's time to design your robot MVP (Minimum Viable Product)! Remember the guidelines you defined.

Where and when

What place?
What time of day?
Does the place or time change?

Draw a picture

What does the robot look like?
Is it attached to something?
Does it move around?
Can its appearance be modified?

Robot's role

Is the robot a friend? Teacher? Helper?
Something else?

Personality

Does the robot have specific characteristics?
Does it have emotional states, or needs?

Context-based behaviour

What external and environmental factors affect behaviour?
What data is used to adapt to context?

Connection to systems

Is the robot connected to external systems, such as software, databases, or other robots?
How does it use these systems?

Interaction modalities

What modalities are inputs to the robot? What modalities does the robot output?

INPUT

☐ movement ☐ screens

☐ voice ☐ touch ☐ lights

☐ sounds ☐ smell ☐ other _____

☐ gestures ☐ facial expressions _____

OUTPUT

☐ movement ☐ screens

☐ voice ☐ touch ☐ lights

☐ sounds ☐ smell ☐ other _____

☐ gestures ☐ facial expressions _____

Interaction flow

Describe the most important interaction of the robot.
Note: only fill the bottom row if your robot is teleoperated.

Social Robot Co-Design Canvases free version by Minja Axelsson is licensed under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) license. Sponsored by Futureice.
 The Social Robot Co-Design Canvases can be found at <https://osf.io/jg2t8/>
 Cite as: Minja Axelsson. 2020. Social Robot Co-Design Canvases. <https://doi.org/10.17605/OSF.IO/JG2T8>

Fig. 10. The MVP canvas, which combines elements from canvases (05) Environment, (06) Form, (07) Interaction, and (08) Behaviour.

4.4.1 (04) MVP Canvas. The MVP canvas, pictured in Figure 10, was constructed to be a summary of canvases (05) Environment, (06) Form, (07) Interaction, and (08) Behaviour. As discussed in Section 3.2, the canvas was created due to requests from participants for a solution space canvas that can be used when the design team has limited time. It includes the following features from canvases (05)–(08): (05) Environment—where and when, connection to systems; (06) Form—draw a picture; (07) Interaction—interaction modalities, interaction flow; (08) Behaviour—role, personality, and context-based behaviour. The features were selected based on their perceived importance to the participants.

4.4.2 (09) Service Ecosystem and (10) Experience Flow Canvases. Two additional canvases, (09) Service Ecosystem and (10) Experience Flow, were created after the canvases were used in two projects. Roboticists reviewing the canvases requested more tools to analyze how the robot fits within its operation context. These canvases examine the relationship of the designed HRI to the broader service.

Service ecosystem canvases have been previously used to describe the relationships present within and around a service (e.g., the *platform canvas*, which emphasizes the central role of

Fig. 11. Service ecosystem canvas.

core interaction toward value capture and monetization [54]). The canvas (09), pictured in Figure 11, examines the relationship between the robot or robots, external technical systems (databases, sensors, actuators, and others), primary and secondary users, and other stakeholders. These different stakeholders can interact with each other via data, information, value, money and credits, or goods and services. The designers are encouraged to draw the relationships between these entities with different coloured arrows—each indicating the resource being exchanged.

The Experience flow canvas (10), pictured in Figure 12, expands on the *interaction flow* factor of the Interaction canvas (pictured in Figure 6). It examines the experience of interacting with the robot in more detail—how the user feels, what they think, and what they do—before, during, and after using the robot. The robot's side of the interaction is described by what it is doing, its sensor input, and its connection to other systems. Additionally, a robot operator's actions can be detailed in the experience flow if the application involves teleoperation. Approaching the whole experience in more detail aids the designers in empathizing with the user throughout the user path, as well as thinking about what operations the robot should be executing, and what data it should be receiving or storing at each point.

Fig. 12. Experience flow canvas.

5 EVALUATION OF THE FRAMEWORK AS A TOOL IN PARTICIPATORY DESIGN

We present the evaluation of the canvases, based on data gathered from 2 projects that employed the canvases (01)–(08). In the first project, roboticists created a gaming robot together with Yle, Finland’s national media broadcast company. The aim of the robot was to play games with teenagers online. In the second project, roboticists created a robot for the Helsinki central public library Oodi. The aim of the robot was to guide customers to books in the library. This section details the workshopping methods used in the projects, what design guidelines and ethical considerations the design teams created, and how the robots performed in their tasks.

5.1 Methods of Evaluation

Both projects followed similar design processes, depicted in Table 1. First, domain specialists were interviewed about the robot’s future operation environment, their daily tasks, and the domain experts’ thoughts about robots’ potential in the environment. Interviews have been previously used to gather initial information in PD for HRI [6, 57, 64].

Next, brainstorming was done based on the information gathered from the interviews. Participants were encouraged to create as many ideas as they could. Ideation has previously been used

Table 1. Project Phases, where Project 1 (P1) is the Gaming Robot Project with Finnish National Public Broadcasting Company and Project 2 (P2) is the Library Robot Project with Helsinki Central Library

1. Interviews	2. Brainstorming	3. Workshop 1 Canvases (01)–(04)	4. Workshop 2 Canvases (05)–(08)	5. Pilot study
P1 5 media production specialists	49 ideas	Workshop 1.1: 4 roboticists (R01, R02, R03, R04); 5 media production specialists (E01, E02, E03, E04); 1 librarian (E05)	Workshop 1.2: 2 roboticists (R01, R04)	Teleoperated Furhat robot. 6 hour pilot study online on Minecraft and Twitch. 431 unique viewers, 49 simultaneous viewers. 16 players responded to surveys afterward.
P2 4 librarians	73 ideas	Workshop 2.1: 3 roboticists (R02, R03, R04); 6 librarians (E05, E06, E07, E08); 3 users (U01, U02, U03)	Workshop 2.2: 3 roboticists (R01, R03, R04)	Autonomous MiR200 robot with built-on module. 6 hour pilot study in Oodi library. 95 users initiated missions, of which 88 were completed successfully.

in HRI workshops [64]. Of these ideas, three were selected to be refined in the first workshops. The first workshop had participants split into three teams. The canvases (01) Problem space, (02) Ethical considerations, (03) Design Guidelines, and (04) MVP were used by each team. At the end of the workshop, the teams presented their ideas and selected which idea should be implemented.

In both design processes, a second workshop was then held to refine the design of the robot, based on what was filled in the (04) MVP canvas in the previous workshop. The filled in canvases from the first workshops were displayed in the room during the second workshops, to make the information readily available. Unfortunately, the full design teams were not able to attend these second workshops and were instead attended only by the roboticists available. The second workshops used the canvases (05) Environment, (06) Form, (07) Interaction, and (08) Behaviour. Based on the information gathered in these workshops, implementations of the robots were created by the same roboticists who were present in the workshops. The implementations are discussed further in Sections 5.2 and 5.3. The feedback of the participants was analyzed for themes of advantages and other noteworthy observations, which are presented in Section 6.

5.2 Gaming Robot Project

The gaming robot project was completed with Finland's national public broadcasting company, Yle. Attendees of Workshop 1.1 were 4 roboticists (R01, R02, R03, and R04)—three of whom acted as facilitators within their respective groups—and five domain experts of media production from Yle (E01, E02, E03, and E04, one of the experts did not provide answers to the survey). Additionally, a librarian and a robotics expert (E05) from the second project attended for cross-pollination of information. Of 49 created ideas, 3 were selected for further workshoping: a robot to play games

online with teenagers, a robot to discuss in public transport with people and provide networking opportunities between nations, and a beggar robot that gives information about begging. Based on what would be most useful to the media broadcast company, the participants had a vote and selected the first option.

5.2.1 Design Guidelines and Ethical Considerations. During Workshop 1.1., the design team created the following design guidelines for the gaming robot:

Environment—The application of the robot should take into account the existing online culture and rules and work to create its own within this environment.

Form—The robot is highly reactive: It reacts with voice, songs, facial expressions, and sound effects.

Interaction—The interaction should be explorative, and an experience for the user. The user should feel surprised. The robot leads the experience but takes input from users.

Behaviour—The robot is clearly a robot but uses humanlike features such as humour and context understanding in communication. It plays the game with human rules and plays on a human's speed.

Based on these guidelines, the team created the character "IQ-201" for the robot. Its character was based on aggressive online gamers, that "are convinced of their own superior intelligence." The robot needed to be rude and reactive to elicit reactions from its teenage audience. To achieve the rudeness, a "rage-meter" was created for the robot. This encouraged gamification of interaction with the robot.

Before implementation, the team also wanted to take into account the following ethical considerations:

Physical safety—The potential overheating of the robot's motors should be monitored by a human.

Data security—The application should be checked for **General Data Protection Regulation (GDPR)** congruency.

Transparency—The application should be honest toward the human about how it works and who is controlling the robot.

Equality across users—The robot should be gender neutral, so that the content it creates is inviting and suitable for all genders.

Emotional consideration—The robot should be equally rude toward everyone and react equally to all hate. However, it should not be too rude so as not to hurt the users' feelings. The person who controls the robot needs to have some rules.

Behaviour enforcement—The robot has to be rude, but there is a line: We do not want users to copy its behaviour. It should be evaluated and modified according to feedback. Additionally, the chat of the Twitch application should be moderated, so no bad behaviour is tolerated.

5.2.2 Robot Implementation. The MVP of the robot's design, based on the aforementioned design guidelines and ethical considerations, can be seen in Figure 13. The robot was implemented over 2 months. Based on the filled in canvases in both design workshops, technical specifications for the robot were created by the roboticists. Behaviour documentation for this robot was created together by the roboticists and media professionals who were present during the workshops, by referring to the canvases and guidelines determined during the workshops. Additionally, media professionals created the online gaming environment for the robot in Minecraft. Given the set-up and nature of the study, we unfortunately did not have access to the prospective users of the robot, i.e., the teenagers. We, however, argue that the canvases can be used by marginalized stakeholders, i.e., users who lack the technical skills to design robots and are often excluded from the design process (such as, the teenagers of this study). The integration of these marginalized stakeholders

Fig. 13. One of the three design teams from Workshop 1.1 details a design for the gaming robot, which was chosen for implementation. The MVP canvas depicts a preliminary plan for the implementation of the robot.

is made possible by the canvases' language being selected to be suitable for people with no previous experience with robotics, and by the interface to the design tool itself, requiring the users to simply give their ideas and feedback by adding post-it notes.

To fulfill all requirements defined by the design team, we selected the Furhat robot. The team decided that to meet the complex personality requirements of the robot, it should be teleoperated. An interaction script for the robot's utterances was created together with the media production experts. The Furhat robot had a relatively easy-to-use teleoperation interface, which allowed for the teleoperator to input text to be turned into the robot's speech, as well as perform gestures at the click of a button.

A pilot study with the robot was performed, where it streamed its gaming online on the platform Twitch, and played the game Minecraft in multiplayer mode on the media company's server. The robot can be seen in Figure 14. The pilot reached 431 unique viewers, with 49 simultaneously

Fig. 14. The gaming robot, created together with Yle (Finland’s national broadcasting company).

online. 16 people responded to a survey about the pilot afterward. According to the survey, 80% of players were under 18, and most were 13- to 15-year-olds; 80% of players interacted with the robot; and 75% of players rated the robot as 3 or above of 5. This indicated that the robot—which was designed with the canvases—successfully met its intended purpose. One teenager commented that the robot was a bit too rude and induced anxiety in them. The design team decided that the robot should be modified to be a bit more friendly in future interactions. One teenager remarked on the robot “looking a bit blue in the face.” This indicates that the robot’s design should be modified further to avoid the uncanny valley effect [66]. The robot was also described as “fun and cool,” and several teenagers remarked they were quite sure it was operated by a human.

The teenagers’ behaviour during the interaction was regarded as interesting by the media production experts. The robot’s provocative character provoked teenagers into repeatedly killing it within the Minecraft game platform. After continued harassment of the robot, we modified the robot’s behaviour script to be more friendly, to create more constructive interactions. There were two clear factions in the game: Some teenagers were intent on killing the robot throughout the entire game, and some protected it throughout. Some became more comfortable with the robot as the stream went on, switching sides. Some teenagers aimed to “make friends” with the robot from the start, giving it flowers and even reassuring it when it was killed within the game. Discussions about the robot’s mode of operation were had throughout the game (i.e., whether it was autonomous or there was someone operating it). The design team decided that for future implementations, it should be made more clear that the robot was following an interaction script.

5.3 Library Robot Project

The library robot project was completed with Helsinki’s central library Oodi. As seen in Table 1, attendees of Workshop 2.1 were three roboticists (R02, R03, and R04) who acted as facilitators within their respective design groups, six librarians from Oodi (E05, E06, E07, and E08, two librarians did not provide answers to the surveys), and three library customers (U01, U02, and U03).

Of 73 created ideas, 3 were selected for further workshoping: the robot to guide customers in the library to books, a robot to help and provide information in multiple languages, and a robot to book workspaces within the library and guide users to them. Based on what would be most useful to the library, the participants had a vote and selected the first option.

5.3.1 Design Guidelines and Ethical Considerations. Design guidelines defined by the team in Workshop 2.1 were as follows:

Environment—The robot’s implementation should allow accessibility despite physical limitations. It should take into account the library’s level of background noise, its changing floor plan and furniture, as well as the moving customers inside the library.

Form—The robot should have a touch screen for interaction. Mobility would be an advantage, however taking users to shelves is not strictly necessary. RFID could be used in the future to find the exact locations of books.

Interaction—Interaction is done via touch screen. It should be noted that the robot was first intended to speak, however this feature was removed so that users would not get the idea that they could speak to the robot. Instead, the robot was decided to beep expressively, and be sincerely a machine in “personality.” Initially, the robot was planned to print a map or give a QR code to the user for them to locate the book, however, when the availability of the MiR200 mobile robot was confirmed, this feature was removed and guiding physically to the book was instead implemented.

Behaviour—The robot should treat users equally.

The design team additionally mentioned monetary savings (i.e., saving the librarians’ time) as a design guideline for the robot’s advantages. The design team also decided that the robot should not be too humanoid. An abstract form for the robot was desirable, with expressive, non-speaking forms of communication.

The team also wanted to make sure that the robot aligned with Oodi’s strategy and policies. The following ethical considerations, which can also be seen in Figure 15, were underlined:

Physical safety—Children can be rough with the robot, there should be a staff member ready to intervene nearby. The robot should not drive on top of people: This can be achieved with the laser sensors integrated in the MiR200 robot, stopping it when obstacles are nearby.

Data security—The application of the robot should be ensured to be GDPR compliant. Specifically, data on the users of the robot should not be combined with what books or book categories they search for.

Transparency—Not considered during the workshop due to lack of time. Roboticists later thought it important that customers should have information on the operation principles of the robot available.

Equality across users—All users should be able to use the robot, regardless of previous experience with technology. The robot should be accessible despite physical restrictions. As the robot is taken into use, new languages should be integrated. Additionally, the robot should use as gender neutral communication as possible (which in the Finnish language is possible by using the gender neutral pronoun “hän,” instead of gendered pronouns such as “she” or “he”).

Emotional consideration—Customers may get frustrated if the robot does not work as intended, and human staff need to be prepared for this. Users have a conflict of interest: They want swift service but also appreciate in-depth service and additional information that they receive from human staff. In time, they will learn this distinction (between the robot and human staff). Users will also need to learn to trust the robot, this can be helped with design choices.

Behaviour enforcement—If users get used to the robot’s speed, then they might become rude to staff. Staff should have a policy in place for this situation. An unresolved design consideration

Fig. 15. One of the three design teams in Workshop 2.1 designed a robot that would help customers in multiple languages. This robot was not selected for implementation. Here, the team details ethical considerations related to the design.

was whether users may transfer rude behaviour from the robot (if it does not protest), toward humans. This will be measured for during the pilot.

5.3.2 Robot Implementation. The robot was implemented over 4 months. Information gathered with the canvases during the workshops guided the roboticists (who also participated in the workshops) and the library personnel in drafting the technical specifications for the robot. The mobile robot MiR200 was selected for realization, due to its suitability and adaptability to this use case. Behaviour documentation for the robot was created and implemented by the roboticists. The robot's behaviour was continuously tested in the library environment, receiving feedback from both the library personnel involved in the workshops and library visitors (spontaneous feedback). During

Fig. 16. The library robot, created together with Oodi public library in Helsinki.

initial testing, the roboticists noticed that the mobile robot by itself was too abstract to form a social bond with people. Children climbed on top of it, and adults did not regard it as a medium for interaction. It was decided that the robot needed a kind of “face” to indicate its capability for social interaction. For this purpose, mechanical “googly eyes” for the robot were constructed. This is pictured in Figure 16.

The robot was piloted autonomously at the library for 6.5 consecutive hours, during which it successfully performed 88 missions, of 95 missions scheduled. This indicated that the robot—which was designed with the canvases—successfully met its intended purpose. Library customer comments were collected by the roboticists during the customers’ interactions with the robot at the library. The abstract form of the robot received positive feedback on its “cuteness,” and its eyes. The fact that it was abstract and not humanoid seemed to make it more approachable. “It makes nice sounds, and doesn’t feel dangerous. I like it more because it is simple and not like a human,” was one older woman’s comment. “I am surprised by how well it works,” commented a young man. On an interesting note, one young woman was initially surprised by the robot moving around behind her, and cursed at it. The woman later returned to apologize to the robot. This indicates that the robot achieved some level of “lifelikeness” through its design. However, for future versions of the robot, a more sophisticated sound signalling system for moving around should be developed, so that users are not surprised.

6 RESULTS

Data about the usefulness of the canvases was collected in two ways during both of the design processes. Facilitators of the first workshops (1.1 and 2.1) took notes of the perceived strengths and weaknesses of the canvases, which were employed to make improvements to the canvases after the workshops. Additionally, all participants were asked to answer surveys with questions after the

workshops. Fifteen of 18 unique participants responded. Three roboticists and one domain expert took part in both projects—the evaluations of these individuals were combined to one evaluation so as not to multiply results. Of these 15 unique participants, 9 were women and 6 men, between the ages of 20 to 60 (due to anonymization, more specific data on ages cannot be presented).

Open questions were used so as not to introduce false dichotomies into the results. Specific questions asked were as follows: “How did using the canvases feel?,” “What was good?,” “Was something unclear?,” “Was something missing?,” “What could be improved?,” “Specific comments about canvases Problem space, Design guidelines, Ethical considerations, or MVP?,” and “Other comments?” We conducted a qualitative thematic analysis of the data provided by the participants, as well as of the facilitators’ notes taken during the PD workshops, to determine the advantages and disadvantages of the canvas tools. To conduct our analysis, we searched the data for themes and patterns, following the process outlined by Braun and Clarke [14]. The analysis was performed by reviewing all 15 participants’ feedback statements for comments about the usefulness and points of improvement of the canvases, and then binning these comments into potential themes. The three advantage themes found by such analysis are presented in Section 6.1 (e.g., for the advantage “structure and clarity of the process,” statements that were considered to indicate this were “made thinking clearer” and “helped in constructing thoughts into ideas”). Additionally, other advantages and noteworthy observations that were present in the data, but that did not constitute major themes, are discussed in Section 6.2. Extracts from the participants’ submitted data, as well as from the facilitators’ notes, are presented as identifiable and representative examples of each theme as part of our analysis.

6.1 Advantages Indicated by Feedback

Advantage themes found were as follows: (1) structure and clarity of design process, (2) sharing multiple perspectives and converging toward a shared viewpoint, and (3) workshoping with the canvases is an effective and pleasant method of design, which helps the participants learn about social robots. All 15 respondents indicated at least one of the three aforementioned advantage themes. These three advantages were generally indicated in response to the question “What was good?,” although some statements related to these advantages were also made in response to the questions “How did using the canvases feel?,” “Specific comments about canvases Problem space, Design guidelines, Ethical considerations, or MVP?,” and “Other comments?”

6.1.1 Structure, Process, and Clarity. Ten of 15 unique respondents indicated benefits of the framework related to structure, process and clarity. Three participants independently reported that the canvases had made their thinking clearer (R01, R03, and E08). Three reported the “canvas was good as a tool” (R01, R03, and E03). Six indicated that the canvases had a clear process or structure for designing a social robot (R01, R02, E03, E05, E06, and E07). Four said they were exhaustive in reviewing the subject of social robot design (R04, E06, U01, and U03).

The clear progress from problem definition to solution definition was also noted by facilitators. Initially, discussion revolved around who the robot was for, and what problem it was trying to solve. During the final canvas (04, MVP), teams were more focused on feature definition, such as how the robot should interact verbally. For example, in Workshop 1.1 team member E02 remarked (as paraphrased by the note-taking facilitator) “This could be a never-ending discussion [between the robot and its user], how does the conversation stop, or how is it stopped?,” to which E04 responded: “The robot could suggest stopping the conversation at certain intervals.”

6.1.2 From Multiple Perspectives to Shared Viewpoint. Eight of 15 unique respondents indicated advantages related to sharing multiple perspectives, and converging toward a shared viewpoint. Six participants said they enjoyed that everyone could share their thoughts (R02, E03, E04, E07,

E08, and U02). Four enjoyed the new, versatile viewpoints (R02, E06, E07, and U02). Three enjoyed the cross-disciplinary collaboration (E02, E03, and U02).

This result is also supported by the observations of the facilitators of the design teams: They noted that while some participants needed a bit of encouragement, everyone eventually took part in ideating and placing post-its on the canvases.

6.1.3 Canvases and Workshopping Pleasant Format That Encourages Learning. Nine of 15 unique respondents enjoyed the way the canvases and workshop were organized. Seven said it was pleasant (R02, R03, E01, E02, E03, E08, and U03). Five said it was interesting (E05, E06, E07, E08, and U01). Two domain experts explicitly mentioned it being a “learning experience” (E04 and E06).

One facilitator remarked that the teams in workshop 2.1 were clearly very excited about the fact that the robot would be built and deployed in the library. The library’s customers who were present in the workshop wanted to come and use the robot once it was finished. This reflects back on one of the advantages of PD introduced in Section 2: Experts of a system and users of a system are often most affected by changes within the system, and their inclusion in PD can be beneficial.

6.2 Other Advantages and Noteworthy Observations

Additional advantages that were not as explicitly indicated by the feedback, but that are argued for by the authors by relying on the theoretical basis presented earlier, are discussed here. Additionally, other noteworthy observations on the use of the framework are discussed. Observations in the feedback were indicated in response to the questions “How did using the canvases feel?,” “Was something unclear?,” “Was something missing?,” “What could be improved?,” “Specific comments about canvases Problem space, Design guidelines, Ethical considerations, or MVP?,” and “Other comments?”

Flexible tool for different contexts—One of the main advantages of the tool is that it can be used in different domains within HRI. While this advantage was not explicitly indicated by user feedback, it can be seen that it has been successfully indicated due to the successful completion of two social robot projects with it. These robots were deployed in two different contexts, and developed with two different types of domain experts (one also involving users). As discussed in Section 2, our design framework distinguishes itself from previous tools by being a broad tool for the design of social robots, which is adaptable to different contexts.

Ethical considerations—A main advantage of the framework is the explicit implementation of ethical considerations within the design process. Roboticists R02, R03, and R04 all considered the ethical considerations presented by the framework as especially important in their feedback. R03 especially remarked that “we had a great discussion about ethical considerations” during Workshop 2.1. Additionally, one participant who did not respond to the survey remarked excitedly during Workshop 1.1 “these [ethical considerations] are great, these are broad.” While this advantage was not as clearly indicated in the feedback as the aforementioned three advantages, it is still considered significant by the authors, as our framework distinguishes itself from other design frameworks in the field of HRI by explicitly considering ethics, as discussed in Section 2.

Facilitator aids process—Previous HRI workshops have also used facilitators [64] to provide explanations when needed, and to categorize the ideas of participants. In our workshops, 6 of 15 participants independently remarked about the benefit of having a facilitator for the workshop sessions. This indicates that the canvases are best used with the help of a facilitator who has familiarized themselves with the tool beforehand. For example, three workshopers (R04, U01, and U02) noted that the facilitators in each team had an important role in helping the team get started by going through the parts of the canvas, and explaining the terms when asked. R04—who was facilitating their team—remarked that the “question prompts written on the canvases are helpful,

but might be ignored during the excitement of designing.” The facilitator has an important role in making sure the team works together and uses the canvases as intended. However, it should be emphasized that the canvas is only a tool: The users are free to adapt it to their needs. As U03 remarked it is “important for the facilitator to note that the idea is not to debate within the team [of the correct answers or correct ways to use the canvases],” rather it is to exchange ideas and build together.

Modifiability—One facilitator remarked that during Workshop 1.1, not all dimensions of canvas (01) problem space were relevant for the team, since their robot would be operating in a digital space. Due to the facilitator’s expertise, the team was able to move forward without getting stuck on these factors. As evident in Figure 13, this design team had modified the canvas to suit their needs, crossing out interaction modalities they did not need, and circling a thing they found especially important with a “YES!!” This is how the canvases are intended to be used: They are a guide for design and should be adapted to each team as needed. However, the modifications are more likely to succeed in the presence of a facilitator who is familiar with the tool beforehand, and is aware of how it would be useful to adapt it to the robot being designed.

Context of use—The canvases were used in a commercial environment in our two case studies, where the canvases worked efficiently as a tool. However, it should be noted that other contexts may work differently. To verify the usefulness of the canvases in an academic context, future research is needed. Additionally, if users are to be included in the design process, then it should be considered whether the information provided in the canvases is presented on a helpful level. During our case studies, R02 and R03 both remarked that some time was used to explain the concepts introduced on the canvases. In the case of different types of users and experts, the language of the canvases may be considered too difficult, or too simple. For example, a more simple version of the canvases could be created for children if they are the intended user group and are to be included as co-designers. An additional contextual consideration is the culture in which the canvases are used. Both of the case studies were performed in Finland. PD has its roots in Scandinavia, and the practice of PD may vary by region [84]. The applicability of the tool in different cultures is a question worthy of further examination.

7 DISCUSSION

7.1 PD Contributions

The features of successful PD were discussed in Section 2. We argue that the major concerns of PD discussed in that section, which we will italicize here, were indeed present in the two presented design processes [98]. Our design processes *take into account the expertise of specialists* and thus includes *multiple viewpoints*, by including specialists operating in the robots’ contexts into the design process, and by allowing them to share their knowledge with the shared language provided by the design tool. In addition, we allow each participant an equal social platform within their respective design teams to share their opinions, as well as their *authentic experiences* of the field and *context* that the robot would operate in. Finally, we demonstrate we have developed a *hands-on method*, as two concrete applications were addressed by designing with the proposed tool.

Spinuzzi et al. [89] argue that although participatory design lacks coherent evaluative criteria, a list developed by them can be used as nascent criteria for internal integrity. We argue that the proposed methodology meets these criteria to a satisfactory extent. Spinuzzi et al. outline the following criteria:

Quality of life for workers—Both projects focused on creating social robot applications proposed by the domain experts themselves (the media professionals for the gaming robot of Section 5.2 and librarians and customers for the library robot of Section 5.3) through a reflective

decision-making process. The proposed canvases allowed for a level of engagement of such workers that would be non-trivial to achieve otherwise.

Collaborative development—Mechanisms for consensus and agreement were demonstrated throughout the workshops with the canvas tool, where design teams negotiated their decision making among themselves. Common language for robotics terms was also provided by the canvas tool. Common aims were defined by the design teams themselves, by formulating design guidelines and ethical considerations to summarize what to prioritize during the robots' development.

Iterative process—After workshops with the canvas tools, domain experts were actively involved in the iterative process of creating the robots. In the case of the gaming robot, media professionals led the development of the robot's behaviour documentation, interaction script and "personality." Users of the robot were then asked after initial deployment of the robot to give their opinions. Ideally, users (i.e., the teenagers interacting with the final version of the robot) could have joined earlier in the design process, but were unfortunately unavailable due to project limitations. In the case of the library robot, librarians were involved in giving feedback on the robot's development throughout. It was continually deployed and tested in the library environment during its development, and both librarians and library customers volunteered their feedback to the roboticists working on the robot.

Additionally, we consider that the implementation of two robots (which successfully achieved the expected results, as discussed in Sections 5.2.2 and 5.3.2) facilitated by the use of the canvas tool is an indirect measure of the usefulness of the tool to facilitate PD.

7.2 The Tool in the Context of RtD

As discussed in Section 3.2, Zimmerman et al. define the criteria for the evaluation of the quality of contributions realized through the RtD approach as *process*, *invention*, *relevance*, and *extensibility* [108]. We argue that our canvas tool, created with the RtD approach to function as a tool to facilitate PD, meets these criteria.

Process—We have documented in detail both our process of creating the canvas tool, as well as the process of applying it to two robot design projects. We have introduced how the need for such a framework initially arose in a robot design project [5], and how the framework's dimensions were selected, tested, and refined in multiple iterations (see Section 3.2). Similarly, we have documented two examples of robots designed with the PD tool, a process that can be reproduced with the canvas tool presented in this article.

Invention—We demonstrate that in the canvas tool, we have integrated various subject matters to create a novel invention. As discussed in Section 2, there are no previous flexible context tangible tools for the PD of social robots. We combine knowledge from the areas of design and robotics, as well as the knowledge of the participants who took part in the evaluation of the framework throughout its iterations, to create the canvas tool presented here.

Relevance—We argue that our canvas tool is relevant to the HRI community, as it enables multidisciplinary teams to use a flexible context, structured and tangible tool for the PD of social robots, which was not previously available. The framework enables multidisciplinary teams (a common aspect in HRI design teams) to consider multiple viewpoints while solving problems by introducing a shared language.

Extensibility—We introduce our canvas tool to enable the HRI community to apply it to their PD of social robots. We encourage future researchers to leverage the knowledge provided by the tool and the two example projects and make modifications to the canvas tools and the design processes as they see fit.

7.3 Ethical Contributions

Our presented design framework encourages the *reflective practice of design*, an important aspect of PD [98], by explicitly introducing ethical considerations and devoting a significant portion of the design processes to their deliberation. As discussed in detail in Section 4.1.2, participants were encouraged to actively anticipate potential future ethical issues, and how they might respond to them in advance. We did this by providing a list of six ethical considerations often cited in HRI literature, and providing participants a framework of consideration: how each consideration may present as a problem on the user's or robot's side of the interaction, and how each problem could be solved or at least mitigated on the user's or robot's side.

Technology philosopher Verbeek argues that technology plays an important role in the coming about of human actions, and as such, it raises the need to find concrete, material answers to the ethical question of "How to act?" [103]. This is a valuable insight, as it assigns the role of a moral actor to the designers and creators of such technologies. Since technology can guide and influence human moral behaviour, designers should take steps to consider what type of moral behaviour the technology they are designing will encourage or induce. In HRI, ethical considerations often exist between the robot and its user [104], with interactions between the two carrying moral weight. However, by taking into account Verbeek's view that engineers and designers are in fact influencing what types of interactions can exist between a technology and its user, we can conclude that the designer of the robot is also a moral actor within the system.

With our canvas tool, we aim to introduce a structured way to start considering the designer's role as a moral one. While our list of considerations is not exhaustive, we consider it a valuable contribution to the field of HRI, as it introduces a method of consciously anticipating potential adverse effects ahead of creating an implementation of a robot. Such tools have not been previously available. However, this process is inherently complex as the behaviour of technology in novel contexts cannot always be anticipated, and carries with it uncertainty. As such, we argue that the consideration of the ethical outcomes that result from the introduction of robots in real world contexts, should be an iterative process that spans from the early design and implementation stages and continues throughout the period during which the robot remains active. This allows a continuous actualization and adjustment of the robots' behavior that takes into account ethical and environmental factors that might have been unforeseeable in the early design stages, by reviewing what desirable and undesirable consequences the robot is triggering within its operating context. This future work could potentially be done by utilizing the canvas (02) Ethical considerations again after the robot has been observed in its application environment, and filling in the canvas again based on new observations. After this, the design team could make further decisions by utilizing other canvases that they would deem relevant to the considerations.

8 CONCLUSION AND FUTURE WORK

There is a lack of a flexible context tool for the PD of social robots. In this work, we worked toward filling such a gap by introducing a design framework as a tool for the PD of social robots, which can be applied by multidisciplinary teams in varied contexts. The tool was developed over 7 design iterations, and received feedback from 97 people. The feedback of 15 people—who used the tool to design two robots—is presented. In general, the main advantages of this tool as identified by the participants include (1) providing structure and clarity during the design process, (2) encouraging teams to work toward a shared viewpoint from multiple perspectives, and (3) being a format that encourages learning and sharing of ideas. In comparison to already existing tools and literature, the canvases proposed in this work are distinctive not only due to their plasticity and flexibility to different application contexts, but also due to their integration of ethical considerations during the design process.

The proposed canvases are extensive, but not exhaustive. In future work, other ethical considerations (such as responsibility [29]) and design variables (such as bystanders in the robot's environment [83, 100]) should be considered and added to the canvases. In addition, future research should also compare the canvases to other available design tools, such that quantitative measures of their effectiveness can be obtained. While the framework introduced here has been shown to achieve the aforementioned advantages, it is not clear what the effects of developing and testing a robot under this framework are, compared to other methods for developing robots. Additionally, research on the canvases' effectiveness in different cultural and organizational contexts, as well as different populations (e.g., children, or users with disabilities) is also needed. Nonetheless, we argue that the proposed framework provides valuable easy-to-use tools to aid researchers, developers and designers in the co-design of social robots.

APPENDIX

A ONLINE RESOURCES

The canvas tool is available with a CC BY-SA 4.0 [25] license in Reference [4] with detailed instructions on how to cite the framework when building on top of it.

ACKNOWLEDGMENTS

We thank all users of the framework for their participation and feedback. We thank Daryl Weir, Annina Antinranta, and Mia Muurimäki for design-related discussions. We thank Sanna Huttunen and Romeo Pulli from Oodi Helsinki public library, as well as Wesa Aapro and Anssi Komulainen from Yle media, for enabling the projects. We thank the Futurice robotics team for aiding with the robot building and testing. We thank Glen Akins for sharing his mechanical and algorithmic implementation of the googly eyes. We thank Elmira Yadollahi and Patricia Alves-Oliveira for their valuable feedback.

REFERENCES

- [1] Chadia Abras, Diane Maloney-Krichmar, Jenny Preece, et al. 2004. User-centered design. *Encyclopedia of Human-Computer Interaction*, W. Bainbridge (Ed.). Sage Publications, Thousand Oaks, CA 445–456.
- [2] Joshua M. Ackerman, Christopher C. Nocera, and John A. Bargh. 2010. Incidental haptic sensations influence social judgments and decisions. *Science* 328, 5986 (2010), 1712–1715.
- [3] Patricia Alves-Oliveira, Pedro Sequeira, and Ana Paiva. 2016. The role that an educational robot plays. In *Proceedings of the 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'16)*. IEEE, 817–822.
- [4] Minja Axelsson. 2020. Social Robot Co-Design Canvases. <https://doi.org/10.17605/OSF.IO/JG2T8>
- [5] Minja Axelsson, Mattia Racca, Daryl Weir, and Ville Kyrki. 2019. A participatory design process of a robotic tutor of assistive sign language for children with autism. In *Proceedings of the 28th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'19)*. IEEE.
- [6] Shiri Azenkot, Catherine Feng, and Maya Cakmak. 2016. Enabling building service robots to guide blind people a participatory design approach. In *Proceedings of the 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI'16)*. IEEE, 3–10.
- [7] Kim Baraka, Ana Paiva, and Manuela Veloso. 2016. Expressive lights for revealing mobile service robot state. In *Proceedings of the 2nd Iberian Robotics Conference (Robot'15)*. Springer, 107–119.
- [8] Christoph Bartneck and Jodi Forlizzi. 2004. A design-centred framework for social human-robot interaction. In *Proceedings of the 13th IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN'04)*. IEEE, 591–594.
- [9] Jenay M. Beer, Arthur D. Fisk, and Wendy A. Rogers. 2014. Toward a framework for levels of robot autonomy in human-robot interaction. *J. Hum.-Robot Interact.* 3, 2 (2014), 74–99.
- [10] E. T. Bekele, U. Lahiri, A. R. Swanson, J. A. Crittendon, Z. E. Warren, and N. Sarkar. 2013. A step towards developing adaptive robot-mediated intervention architecture (ARIA) for children with autism. *IEEE Trans. Neural Syst. Rehabil. Eng.* 21, 2 (Mar. 2013), 289–299.

- [11] Lykke Brogaard Bertel, Dorte Malig Rasmussen, and Ellen Christiansen. 2013. Robots for real: Developing a participatory design framework for implementing educational robots in real-world learning environments. In *Proceedings of the IFIP Conference on Human-Computer Interaction*. Springer, 437–444.
- [12] Erling Björgvinsson, Pelle Ehn, and Per-Anders Hillgren. 2012. Design things and design thinking: Contemporary participatory design challenges. *Des. Issues* 28, 3 (2012), 101–116.
- [13] Laura Boccanfuso, Sarah Scarborough, Ruth K. Abramson, Alicia V. Hall, Harry H. Wright, and Jason M. O’Kane. 2017. A low-cost socially assistive robot and robot-assisted intervention for children with autism spectrum disorder: Field trials and lessons learned. *Auton. Robots* 41, 3 (2017), 637–655.
- [14] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.* 3, 2 (2006), 77–101.
- [15] Cynthia Breazeal. 2001. Emotive qualities in robot speech. In *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 3. IEEE, 1388–1394.
- [16] Cynthia Breazeal. 2004. Social interactions in HRI: The robot view. *IEEE Trans. Syst. Man. Cybernet. C (Appl. Rev.)* 34, 2 (2004), 181–186.
- [17] Cynthia Breazeal, Kerstin Dautenhahn, and Takayuki Kanda. 2016. Social robotics. In *Springer Handbook of Robotics* (2nd ed.), Bruno Siciliano and Oussama Khatib (Eds.). Springer Publishing Company, Incorporated, 1935–1971.
- [18] Cynthia L. Breazeal. 2004. *Designing Sociable Robots*. MIT Press.
- [19] Rodney Brooks. 1986. A robust layered control system for a mobile robot. *IEEE J. Robot. Autom.* 2, 1 (1986), 14–23.
- [20] Tim Brown. 2009. *Change by Design*. Harper Collins.
- [21] Drazen Brscić, Hiroyuki Kidokoro, Yoshitaka Suehiro, and Takayuki Kanda. 2015. Escaping from children’s abuse of social robots. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 59–66.
- [22] Ryan Calo. 2014. *The Case for a Federal Robotics Commission*. Technical Report. Brookings Institution Center for Technology Innovation. Retrieved from https://www.brookings.edu/wp-content/uploads/2014/09/RoboticsCommissionR2_Calo-1.pdf.
- [23] Anupam Chander. 2017. The racist algorithm. *Mich. L. Rev.* 115, 6 (2017), 1023–1045.
- [24] EunJeong Cheon and Norman Makoto Su. 2016. Integrating roboticist values into a value sensitive design framework for humanoid robots. In *Proceedings of the 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI’16)*. IEEE, 375–382.
- [25] Creative Commons. 2020. CC BY-SA 4.0. Retrieved from <https://creativecommons.org/licenses/by-sa/4.0/>.
- [26] Kate Darling. 2012. Extending legal protection to social robots: The effects of anthropomorphism, empathy, and violent behavior towards robotic objects. In *Proceedings of the We Robot Conference*.
- [27] Kate Darling. 2017. “Who’s johnny?” Anthropomorphic framing in human-robot interaction, integration, and policy. In *Robot Ethics 2.0*, P. Lin, G. Bekey, K. Abney, and R. Jenkins (Eds.). Oxford University Press.
- [28] Kate Darling, Palash Nandy, and Cynthia Breazeal. 2015. Empathic concern and the effect of stories in human-robot interaction. In *Proceedings of the 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN’15)*. IEEE, 770–775.
- [29] Edoardo Datteri. 2013. Predicting the long-term effects of human-robot interaction: A reflection on responsibility in medical robotics. *Sci. Eng. Ethics* 19, 1 (2013), 139–160.
- [30] K. Dautenhahn. 2002. Design spaces and niche spaces of believable social robots. In *Proceedings of the 11th IEEE International Workshop on Robot and Human Interactive Communication*. 192–197.
- [31] Kerstin Dautenhahn, Chrystopher L. Nehaniv, Michael L. Walters, Ben Robins, Hatice Kose-Bagci, N. Assif Mirza, and Mike Blow. 2009. KASPAR—A minimally expressive humanoid robot for human–robot interaction research. *Appl. Bionics Biomech.* 6, 3–4 (2009), 369–397.
- [32] Kerstin Dautenhahn and Iain Werry. 2004. Towards interactive robots in autism therapy: Background, motivation and challenges. *Pragm. Cogn.* 12, 1 (2004), 1–35.
- [33] Gemma Derrick. 2018. Risking groupthink in impact assessment. In *The Evaluators’ Eye*. Springer, 159–197.
- [34] Anca D. Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S. Srinivasa. 2015. Effects of robot motion on human-robot collaboration. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI’15)*. ACM, New York, NY, 51–58. <http://doi.acm.org/10.1145/2696454.2696473>
- [35] Brian R. Duffy and Karolina Zawieska. 2012. Suspension of disbelief in social robotics. In *Proceedings of the 21st IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN’12)*. IEEE, 484–489.
- [36] Gerhard Fischer. 2000. Symmetry of ignorance, social creativity, and meta-design. *Knowl.-Bas. Syst.* 13, 7–8 (2000), 527–537.
- [37] Gerhard Fischer, Kumiyo Nakakoji, and Yunwen Ye. 2009. Metadesign: Guidelines for supporting domain experts in software development. *IEEE Softw.* 26, 5 (2009).

- [38] Gerhard Fischer and Eric Scharff. 2000. Meta-design: Design for designers. In *Proceedings of the 3rd Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*. ACM, 396–405.
- [39] Batya Friedman, Peter H. Kahn, and Alan Borning. 2008. Value sensitive design and information systems. In *The Handbook of Information and Computer Ethics*, 69–101.
- [40] Futurice. 2017. Lean Service Creation. Retrieved from <https://leanservicecreation.com>.
- [41] Nicole Giullian, Daniel Ricks, Alan Atherton, Mark Colton, Michael Goodrich, and Bonnie Brinton. 2010. Detailed requirements for robots in autism therapy. In *Proceedings of the IEEE International Conference on Systems Man and Cybernetics (SMC'10)*. IEEE, 2595–2602.
- [42] D. Glas, Satoru Satake, Takayuki Kanda, and Norihiro Hagita. 2012. An interaction design framework for social robots. In *Robotics: Science and Systems*, Vol. 7. 89.
- [43] James Richard Goodwin and Alan Winfield. 2008. *A Unified Design Framework for Mobile Robot Systems*. Ph.D. Dissertation. University of the West of England, Bristol.
- [44] Michal Gordon, Eileen Rivera, Edith Ackermann, and Cynthia Breazeal. 2015. Designing a relational social robot toolkit for preschool children to explore computational concepts. In *Proceedings of the 14th International Conference on Interaction Design and Children*. ACM, 355–358.
- [45] Javi F. Gorostiza, Ramon Barber, Alaa M. Khamis, Maria Malfaz, Rakel Pacheco, Rafael Rivas, Ana Corrales, Elena Delgado, and Miguel A. Salichs. 2006. Multimodal human-robot interaction framework for a personal robot. In *Proceedings of the 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'06)*. IEEE, 39–44.
- [46] Pamela J. Hinds, Teresa L. Roberts, and Hank Jones. 2004. Whose job is it anyway? A study of human-robot interaction in a collaborative task. *Hum.-Comput. Interact.* 19, 1 (2004), 151–181.
- [47] Guy Hoffman, Oren Zuckerman, Gilad Hirschberger, Michal Luria, and Tal Shani Sherman. 2015. Design and evaluation of a peripheral robotic conversation companion. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 3–10.
- [48] Claire A. G. J. Huijnen, Monique A. S. Lexis, Rianne Jansens, and Luc P. de Witte. 2017. How to implement robots in interventions for children with autism? A co-creation study involving people with autism, parents and professionals. *J. Autism Dev. Disord.* 47, 10 (2017), 3079–3096.
- [49] Alexandre Joyce and Raymond L. Paquin. 2016. The triple layered business model canvas: A tool to design more sustainable business models. *J. Clean. Prod.* 135 (2016), 1474–1486.
- [50] Peter H. Kahn, Nathan G. Freier, Takayuki Kanda, Hiroshi Ishiguro, Jolina H. Ruckert, Rachel L. Severson, and Shaun K. Kane. 2008. Design patterns for sociality in human-robot interaction. In *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction*. ACM, 97–104.
- [51] Andra Keay. 2011. The naming of robots. In *Proceedings of the Pioneers Workshop at the 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI'11)*.
- [52] Merel Keijsers and Christoph Bartneck. 2018. Mindless robots get bullied. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 205–214.
- [53] Susan Kippax. 2018. A journey to HIV prevention research: From social psychology to social health via multidisciplinary. *J. Health Psychol.* 23, 3 (2018), 442–456.
- [54] Heidi M. E. Korhonen, Kaisa Still, Marko Seppänen, Miika Kumpulainen, Arho Suominen, and Katri Valkokari. 2017. The core interaction of platforms: How startups connect users and producers. *Technol. Innov. Manage. Rev.* 7, 9 (2017), 17–29.
- [55] Hideki Kozima, Marek P. Michalowski, and Cocoro Nakagawa. 2009. Keepon. *Int. J. Soc. Robot.* 1, 1 (2009), 3–18.
- [56] Hideki Kozima and Cocoro Nakagawa. 2006. Interactive robots as facilitators of childrens social development. In *Mobile Robots: Towards New Applications*. Intech.
- [57] Hee Rin Lee, Selma Šabanović, Wan-Ling Chang, David Hakken, Shinichi Nagata, Jen Piatt, and Casey Bennett. 2017. Steps toward participatory design of social robots: Mutual learning with older adults with depression. In *Proceedings of the 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI'17)*. IEEE, 244–253.
- [58] Kwan Min Lee, Wei Peng, Seung-A. Jin, and Chang Yan. 2006. Can robots manifest personality?: An empirical test of personality recognition, social responses, and social presence in human-robot interaction. *J. Commun.* 56, 4 (2006), 754–772.
- [59] Larry J. Leifer and Martin Steinert. 2011. Dancing with ambiguity: Causality behavior, design thinking, and triple-loop-learning. *Inf. Knowl. Syst. Manage.* 10, 1–4 (2011), 151–173.
- [60] Daniel Leyzberg, Samuel Spaulding, and Brian Scassellati. 2014. Personalizing robot tutors to individuals' learning differences. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-robot Interaction*. ACM, 423–430.
- [61] Robert F. Lusch, Stephen L. Vargo, and Matthew O'brien. 2007. Competing through service: Insights from service-dominant logic. *J. Retail.* 83, 1 (2007), 5–18.
- [62] Joseph B. Lyons. 2013. Being transparent about transparency. In *Proceedings of the AAAI Spring Symposium*.

- [63] Noriaki Mitsunaga, Christian Smith, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2006. Robot behavior adaptation for human-robot interaction based on policy gradient reinforcement learning. *J. Robot. Soc. Jpn.* 24, 7 (2006), 820–829.
- [64] Sanika Moharana, Alejandro E. Panduro, Hee Rin Lee, and Laurel D. Riek. 2019. Robots for joy, robots for sorrow: Community based robot design for dementia caregivers. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI'19)*. IEEE, 458–467.
- [65] Dobrila Rancic Moogk. 2012. Minimum viable product and the importance of experimentation in technology startups. *Technol. Innov. Manage. Rev.* 2, 3 (2012).
- [66] Masahiro Mori. 1970. The uncanny valley. *Energy* 7, 4 (1970), 33–35.
- [67] Michael J. Muller. 2003. Participatory design: The third space in HCI. *Hum.-comput. Interact.: Dev. Process* 4235, 2003 (2003), 165–185.
- [68] Enid Mumford. 2000. A socio-technical approach to systems design. *Require. Eng.* 5, 2 (2000), 125–133.
- [69] Bilge Mutlu and Jodi Forlizzi. 2008. Robots in organizations: The role of workflow, social, and environmental factors in human-robot interaction. In *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction*. ACM, 287–294.
- [70] Davide Nicolini, Jeanne Mengis, and Jacky Swan. 2012. Understanding the role of objects in cross-disciplinary collaboration. *Org. Sci.* 23, 3 (2012), 612–629.
- [71] Jakob Nielsen. 1993. Iterative user-interface design. *Computer* 11 (1993), 32–41.
- [72] Kei Okada, Mitsuharu Kojima, Yuichi Sagawa, Toshiyuki Ichino, Kenji Sato, and Masayuki Inaba. 2006. Vision based behavior verification system of humanoid robot for daily environment tasks. In *Proceedings of the 6th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 7–12.
- [73] Cristina A. Pop, Ramona E. Simut, Sebastian Pintea, Jelle Saldien, Alina S. Rusu, Johan Vanderfaeillie, Daniel O. David, Dirk Lefeber, and Bram Vanderborght. 2013. Social robots vs. computer display: Does the way social stories are delivered make a difference for their effectiveness on ASD children? *J. Educ. Comput. Res.* 49, 3 (2013), 381–401.
- [74] Eeva Raita and Risto Sarvas. 2017. Blending lean, design, and agile thinking into one: Lean service creation for facilitating multidisciplinary teams. *Touchpoint* 8, 3 (2017), 74–79.
- [75] Byron Reeves and Clifford Ivar Nass. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press.
- [76] Laurel Riek and Don Howard. 2014. A code of ethics for the human-robot interaction profession. In *Proceedings of the We Robot Conference*.
- [77] Lionel Robert. 2018. Personality in the human robot interaction literature: A review and brief critique. In *Proceedings of the 24th Americas Conference on Information Systems*. 16–18.
- [78] Jennifer Robertson. 2010. Gendering humanoid robots: Robo-sexism in Japan. *Body Soc.* 16, 2 (2010), 1–36.
- [79] Ben Robins, Kerstin Dautenhahn, and Janek Dubowski. 2006. Does appearance matter in the interaction of children with autism with a humanoid robot? *Interact. Stud.* 7, 3 (2006), 479–512.
- [80] Martin Saerbeck and Christoph Bartneck. 2010. Perception of affect elicited by robot motion. In *Proceedings of the 5th ACM/IEEE International Conference on Human-robot Interaction (HRI'10)*. IEEE Press, Piscataway, NJ, 53–60.
- [81] Miguel A. Salichs, Ramon Barber, Alaa M. Khamis, María Malfaz, Javier F. Gorostiza, Rakel Pacheco, Rafael Rivas, Ana Corrales, Elena Delgado, and David Garcia. 2006. Maggie: A robotic platform for human-robot social interaction. In *Proceedings of the IEEE Conference on Robotics, Automation and Mechatronics*. IEEE, 1–7.
- [82] Matthias Scheutz. 2011. *Robot Ethics: The Ethical and Social Implications of Robotics*. The MIT Press.
- [83] Jean Scholtz and Siavosh Bahrami. 2003. Human-robot interaction: Development of an evaluation methodology for the bystander role of interaction. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics. Conference Theme-System Security and Assurance (SMC'03)*, Vol. 4. IEEE, 3212–3217.
- [84] Douglas Schuler and Aki Namioka. 1993. *Participatory Design: Principles and Practices*. CRC Press.
- [85] Alessandra Sciutti, Lars Schillingmann, Oskar Palinko, Yukie Nagai, and Giulio Sandini. 2015. A gaze-contingent dictating robot to study turn-taking. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*. ACM, 137–138.
- [86] Masahiro Shiomi, Kazuhiko Shinozawa, Yoshifumi Nakagawa, Takahiro Miyashita, Toshio Sakamoto, Toshimitsu Terakubo, Hiroshi Ishiguro, and Norihiro Hagita. 2013. Recommendation effects of a social robot for advertisement-use context in a shopping mall. *Int. J. Soc. Robot.* 5, 2 (2013), 251–262.
- [87] Roland Siegwart, Kai O. Arras, Samir Bouabdallah, Daniel Burnier, Gilles Froidevaux, Xavier Greppin, Björn Jensen, Antoine Lorotte, Laetitia Mayor, Mathieu Meisser, et al. 2003. Robox at expo. 02: A large-scale installation of personal robots. *Robot. Auton. Syst.* 42, 3–4 (2003), 203–222.
- [88] Jesper Simonsen and Toni Robertson. 2012. *Routledge International Handbook of Participatory Design*. Routledge.
- [89] Clay Spinuzzi. 2005. The methodology of participatory design. *Techn. Commun.* 52, 2 (2005), 163–174.

- [90] Aaron Steinfeld, Odest Chadwicke Jenkins, and Brian Scassellati. 2009. The oz of wizard: Simulating the human for interaction research. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction*. ACM, 101–108.
- [91] Niya Stoimenova, Lenny Van Onselen, and Rianne Valkenberg. 2015. Four guiding factors for facilitators of multi-disciplinary collaboration. In *Proceedings of the 4th Participatory Innovation Conference*. 51.
- [92] Lucy Suchman. 2002. Located accountabilities in technology production. *Scand. J. Inf. Syst.* 14, 2 (2002), 7.
- [93] Dag Sverre Syrdal, Kerstin Dautenhahn, Kheng Lee Koay, Michael L. Walters, and Wan Ching Ho. 2013. Sharing spaces, sharing lives—The impact of robot mobility on user perception of a home companion robot. In *Proceedings of the International Conference on Social Robotics*. Springer, 321–330.
- [94] Leila Takayama and Caroline Pantofaru. 2009. Influences on proxemic behaviors in human-robot interaction. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5495–5502.
- [95] Onur Tekdas, Volkan Isler, Jong Hyun Lim, and Andreas Terzis. 2009. Using mobile robots to harvest data from sensor fields. *IEEE Wireless Commun.* 16, 1 (2009), 22–28.
- [96] Ashutosh Tewari, James Peabody, Richard Sarle, Guruswami Balakrishnan, Ashok Hemal, Alok Shrivastava, and Mani Menon. 2002. Technique of da vinci robot-assisted anatomic radical prostatectomy. *Urology* 60, 4 (2002), 569–572.
- [97] Andreas Theodorou, Robert H. Wortham, and Joanna J. Bryson. 2016. Why is my robot behaving like that? Designing transparency for real time inspection of autonomous robots. In *Proceedings of the AISB Workshop on Principles of Robotics*. University of Bath.
- [98] Bettina Törpel, Alex Voss, Mark Hartswood, and Rob Procter. 2009. Participatory design: Issues and approaches in dynamic constellations of use, design, and research. In *Configuring User-Designer Relations*. Springer, 13–29.
- [99] Pagona Tsormpatzoudi, Bettina Berendt, and Fanny Coudert. 2015. Privacy by design: from research and policy to practice—the challenge of multi-disciplinarity. In *Annual Privacy Forum*. Springer, 199–212.
- [100] Katherine M. Tsui, Munjal Desai, and Holly A. Yanco. 2010. Considering the bystander’s perspective for indirect human-robot interaction. In *2010 Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 129–130.
- [101] Pinar Uluur, Neziha Akalin, and Hatice Köse. 2015. A new robotic platform for sign language tutoring. *Int. J. Soc. Robot.* 7, 5 (2015), 571–585.
- [102] Bram Vanderborght, Ramona Simut, Jelle Saldien, Cristina Pop, Alina S. Rusu, Sebastian Pintea, Dirk Lefeber, and Daniel O. David. 2012. Using the social robot probio as a social story telling agent for children with ASD. *Interact. Stud.* 13, 3 (2012), 348–372.
- [103] Peter-Paul Verbeek. 2006. Materializing morality: Design ethics and technological mediation. *Sci. Technol. Hum. Values* 31, 3 (2006), 361–380.
- [104] Gianmarco Veruggio and Fiorella Operto. 2008. Roboethics: Social and ethical implications of robotics. In *Springer Handbook of Robotics*, Bruno Siciliano and Oussama Khatib (Eds.). Springer, Berlin, 1499–1524. https://doi.org/10.1007/978-3-540-30301-5_65
- [105] Peter Vink, A. S. Imada, and Klaus J. Zink. 2008. Defining stakeholder involvement in participatory design processes. *Appl. Ergon.* 39, 4 (2008), 519–526.
- [106] Kazuyoshi Wada, Takanori Shibata, Tomoko Saito, and Kazuo Tanie. 2004. Effects of robot-assisted activity for elderly people and nurses at a day service center. *Proc. IEEE* 92, 11 (2004), 1780–1788.
- [107] Hongxing Wei, Xinming Duan, Shiyi Li, Guofeng Tong, and Tianmiao Wang. 2009. A component based design framework for robot software architecture. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 3429–3434.
- [108] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 493–502.

Received August 2020; revised March 2021; accepted May 2021