

The Problem of Ableist Paternalism in Assistive Robotics

Pratyusha Ghosh^{1,*}, Belén Lledo^{1,2,*}, Laurel D. Riek¹

¹Computer Science and Engineering, University of California San Diego, USA

²Institute of Philosophy, Spanish National Research Council, Spain

ABSTRACT

Autonomous robots can influence people's behaviors socially, cognitively, and physically. They introduce new forms of paternalism, i.e., acting on behalf of someone's benefit without the consent or will of the person. There has been substantial critiques by the critical disability studies community about the longstanding persistence of paternalism towards disabled people. However, despite the growing numbers of assistive robots for disabled people, there has yet to be a comprehensive framework exploring the different complex factors shaping robot-mediated paternalism and the problems associated with that. Since the robot design lifecycle is complex and involves several stakeholders, each of whom can have conflicting views about how to benefit the user, avoiding ableist paternalism can be challenging. Here, we present a possible theoretical framework that identifies the ways robot-mediated paternalism can happen. We plan to build on this framework to ultimately develop a checklist that guides the HRI community on how to avoid ableist paternalism in the design of assistive robots for disabled people.

ACM Reference Format:

Pratyusha Ghosh^{1,*}, Belén Lledo^{1,2,*}, Laurel D. Riek¹. 2024. The Problem of Ableist Paternalism in Assistive Robotics. In *HRI 2024 Workshop on Assistive Applications, Accessibility, and Disability Ethics (A3DE)*, 2024, Boulder, CO, USA.

1 INTRODUCTION

Autonomous robots operate in dynamic environments such as homes, hospitals, and schools. While this has brought many benefits, such as reducing people's cognitive and physical load when robots take over repetitive or dangerous tasks, it can also impact people's *personal autonomy* [15, 17]. For example, robots have a tremendous capacity to interfere with people's behaviors and decision-making in their daily activities and relationships [8, 35]. This effect is particularly worthy of ethical inquiry in the case of robots developed to assist disabled people, as this is a community whose agency and autonomy have historically been undermined [23, 32, 36].

In particular, *paternalism* has been a long-standing subject of debate within the critical disability studies (CDS) and bioethics communities. Paternalism is the ethico-philosophical concept that refers to the interference of an individual and/or institution of authority with another person's autonomy against their will, in pursuit of their best interests [14]. While necessary in cases in which a person is unable to make decisions for themselves (e.g. is unconscious), the disabled community is more often than not unnecessarily subjected to paternalism in their daily lives [36]. This is primarily due to ableist assumptions from non-disabled people.

Ableism refers to the "stereotyping, prejudice, discrimination, and social oppression toward people with disabilities" [5]. It leads to misguided perceptions of the needs and interests of disabled people and what they can do for themselves as well as the amount

of supposed "good" the non-disabled person's assistance can be for the disabled person [12]. Paternalism is the consequence of non-disabled people claiming a need to have control over how disabled people are cared for to ensure they have an allegedly "better quality of life" for their "own good". This imposed and presumptive help can be demeaning and dismissive of disabled people's autonomy [36]. In the space of assistive robotics, paternalism can often manifest both overtly and covertly in the overall robot design and in the kinds of interactions it has with a disabled user and/or their care network.

For example, here is a scenario: An older adult, María, has recently had a stroke. As a result, she now has difficulties with memory, vision, and balance, amongst others. She lives alone, and her adult son, Matthew, who lives several hours away, is worried about her as she is at risk of a second stroke and other health complications. However, María refuses to let a nurse visit to help her or move to assisted living. Out of concern, Matthew offers María a social robot to monitor her health and send him alerts about her condition at home. María, however, dislikes this and finds it unnecessary to have a robot at home sending Matthew alerts every time something happens. She would like to maintain her independence and asserts that she can take care of herself alone.

In this scenario, the robot's mere presence would be paternalistic if there is an imposition of "help" in the form of a robot. Additionally, the robot would be behaving paternalistically whenever it alerts Matthew about María's condition without her consent. To avoid being paternalistic, there needs to be shared decision-making about the robot's presence and its actions between Matthew and María. However, what should happen if they cannot reach a consensus? Should Matthew leave his mother alone even if she is at risk of having a serious accident or another stroke? Is Matthew trying to "overprotect" María due to his ableist prejudices about her newly acquired mobility impairments? Does the clinician involved in María's care have a say in this situation? If María does accept the robot freely, but then the robot behaves paternalistically, is that acceptable, and to what degree?

It is clear that the multiple stakeholders and their conflicting interests are complex. Negotiation of these conflicting interests can be further complicated by the diversity and constant fluctuations in the experience of disability. It can be challenging to determine what role, if any, ableist prejudices play in a given paternalistic interaction, particularly when a robot is involved. Thus, there is a need for a comprehensive framework that accounts for all these different considerations when designing robots for disabled people.

Paternalism can manifest in different forms and can involve either directly interfering with someone's actions or decisions, or covertly influencing them to act or decide in a certain way (e.g. nudging). Nudging in AI health apps has been scrutinized by ethicists due to their ability to covertly influence a person to engage in behaviors that are supposedly beneficial for them regardless of

*Both authors contributed equally to this research.

their personal preferences [16, 28, 29, 33]. Kuhler [28] argues that the opacity of these paternalistic influences can be irresponsible as users may have misguided trust in these systems. This may lead to erroneous assumptions about what is or is not beneficial for them and other people around them. This ties in with substantial literature within the AI and robot ethics communities regarding the appropriateness of designing deception into technologies to support the well-being of the users, which could be considered inherently paternalistic [21, 40].

Robots are embodied agents that support a wider variety of interactions and are situated within dynamic, human environments such as homes and hospitals. Hence, there is greater potential for them to act paternalistically.

There has been some discussion on how paternalism can manifest in robots' behavior. On the one hand, there is literature more focused on the interaction between users and robots in complex environments, such as the aforementioned ethical accounts of paternalism [6, 22, 28]. On the other hand, scholars in CDS have discussed how ableism can be present in design practices and result in poor or even harmful assistive technology [23, 32].

However, there has not yet been a clear account of the complex factors that influence paternalism in assistive robotics, specifically in relation to disabled people. In order to clearly identify these factors, there is a need to outline the complex power structures and the professional, familial, and interpersonal relationships that may influence different phases of the robot design lifecycle. It is also important to integrate perspectives from CDS as it centers the lived experiences of disabled people without value judgments about the quality of life [23, 32]. We posit that this inclusion of the CDS perspective has not been sufficiently accomplished in the analysis of paternalism in robotics.

Furthermore, though the HRI community has been increasingly promoting the inclusivity of end-users in robot design processes, there is a lack of clear guidelines that help HRI researchers detect and mitigate this multifaceted risk of paternalism.

Our work seeks to address these gaps by providing a roadmap of the factors shaped by ableism that can affect the robot design lifecycle and result in paternalism towards disabled people. As our work is still ongoing, in this workshop paper, we primarily provide an in-depth justification of the importance of analyzing this problem. Additionally, we explain how we plan to address this problem, i.e. by developing a usable, practice-oriented guide for technology designers and/or researchers to determine when paternalistic design is not appropriate, and how they should go about mitigating it when possible to respect individuals' autonomy. We hope to contribute to the ongoing conversation within the HRI community about the importance of reflecting on robot-mediated paternalism by highlighting the specific issues to consider when designing with and for disabled people.

2 RELATED WORK

2.1 Disability, Paternalism, and Robot Ethics

Ableism can manifest in the form of benevolent or charitable paternalism [36]. While disability is a spectrum and the needs and desires of disabled people are not the same, ableism is a systematic form of discrimination they are all susceptible to in different degrees and

ways. Among other aspects, disabled people can be subjected to prejudices and misguided expectations depending on whether they are perceived as disabled or not [5].

In what follows, we focus on the risk of paternalism disabled people face precisely because they are perceived as disabled, and, to define it, we leave aside the differences among their particular disabilities. We refer to it as *ableist paternalism*: the type of paternalism that, while being benevolent or sympathetic, arises from the misguided assumption that disabled people need to be helped, saved, or tutored just because they are disabled. In other words, ableist paternalism is the one which is justified only on the basis of a disability. This type of paternalism happens in the form of unwanted help, infantilization, or dismissal of their decisions under the assumption that they do not really know what they want or need [36].

As asserted by [22], robot paternalism is not the same as human paternalism. Robots are not able to develop the intentions that underlie human paternalism, but they can nevertheless be the agent of paternalistic actions. For example, robots can interfere socially in user's decision-making by covertly nudging them [6], by explicitly persuading and coaching them [7], or by affecting human-to-human interaction [18]. They can also interfere physically when the user depends on the robot for their mobile autonomy: e.g. a robotic wheelchair has a high capacity to override a user's desires.

Some robots are conceived as "helpers for people with disabilities" [43]. In those cases, one crucial factor is the underlying understanding of disability designers have. If disability is conceived as a lack or defect to be remedied by technology, the resulting robot is more likely to be paternalistic.

In order to acknowledge the different forms in which ableist paternalism can arise in HRI as well as the specificities robots create, we propose the following definition of robot-mediated paternalism. Robot-mediated paternalism happens when a robot is involved in an action (or set of actions) that is intended to benefit the user when the user is either unaware of what the objective is, or directly does not want it. Robot-mediated paternalism can occur in relation to the general functional objective (e.g. the robot is intended to promote healthy eating in the user), and can also occur in relation to the manner in which the robot attempts to provide this benefit (e.g. suggesting healthy recipes, alerting the user when they are failing to comply with the plan, collecting information about routines). We talk of "robot-mediated paternalism" instead of "robot paternalism" because when a robot behaves paternalistically, it should be evaluated in the context of the human-led design process that has resulted in this type of robot behavior.

The specific problem of ableist paternalism in HRI has not been sufficiently acknowledged by the robot ethics literature to date. To do so, we turn to CDS literature and its analysis of technology.

2.2 Situating Technology Design in Critical Disability Studies

Several scholars in CDS have analyzed how the design of technology for disabled people can be biased by ableist assumptions and fail to respond to the needs of the intended users. These biases in technology design can also be a result of following the medical model of disability, which oversimplifies and frames disability as a

result of an individual's impairments or symptoms as opposed to being socially constructed due to inaccessible environments and asymmetrical power structures [23, 30, 32]. Shew [2] introduced the term "technoableism" for cases in which the designers assume that disability is a problem to be fixed by technology.

Researchers have also proposed new ways of designing and implementing technology that can be emancipatory for and respectful of disabled people. Collective Access scholarship and advocacy have been determined to shift the focus towards fixing environments and making them more disability-friendly, as opposed to trying to adapt the individual to an ableist environment [19]. In the same vein, Crip Technoscience is a step towards non-ableist technoscience, based on the premise that disabled people are the experts that can generate a politicized design for their everyday needs [20]. While traditional care theories may conceive it as a top-down relationship in which a knowledgeable professional provides the service of care to a passive patient, disability theory has contributed to a more horizontal idea of care. Community care is based on the solidarity and response to each other's needs among disabled people, who may not need medical care but mutual support for daily life. This idea of care contributes to flattening the hierarchies that are characteristic of the medical model of disability and related to the possibility of ableist paternalism [37].

There has been relevant work within the HCI community on how to embed CDS into the design of assistive technologies and navigate the complex trade-offs that may arise. This is essential in order to design technology that can truly respond to the needs of disabled people as opposed to ableist attempts to "help" them by "fixing" their disability [23, 32, 41].

3 ONGOING WORK

In our ongoing work, we are developing a roadmap for a comprehensive account of the elements that affect ableist paternalism in human-robot interaction. To do so, we first show why paternalism may be a problem difficult to detect, and then we propose a path for contributing to identifying and avoiding it in HRI.

When disability is framed under the medical model, paternalism can be claimed to be the appropriate answer to situations in which the person is deemed to be unable to act in their best interests or make their decisions "correctly". Consequently, someone should supervise them, take care of them, or act on their behalf. However, both CDS [10, 37, 39] and new approaches to care theory [9, 13, 45] advance a different understanding of autonomy and care. Autonomy can be better understood as a continuum, shaped by the social surroundings of the person [1, 31]. Whether disabled people need specialized care or not, good care or assistance would always imply the promotion of autonomy as best as possible [42].

Even with the overall objective of promoting autonomy, good care can indeed imply an overriding of a person's decisions in some situations. For example, caregivers of people with dementia usually face situations in which they may need to intervene in a careegee's behavior for safety reasons [34]. These kinds of dilemmas are always significantly delicate, and there may be various stakeholders involved in navigating them, including caregivers and clinicians

[26, 27]. The case of dementia can be enlightening for understanding the complicated path to finding a balance between good care and autonomy.

However, disability is a much wider spectrum and a plethora of less clear situations can arise when there are multiple stakeholders involved in a decision regarding care or assistance. While many disabled people do not need the direct intervention of any caregivers in their daily lives, the design of a robot always involves various stakeholders with different expertise backgrounds, interests, and power [25]. Furthermore, the robot itself will be a new source of potentially paternalistic actions when the disabled person interacts with it.

Some of the stakeholders involved in the robot design lifecycle include end-users, clinicians, robot designers, and others [38]. The robot design lifecycle can be extensive and is split into multiple phases: design, development, deployment, and exit. Across these different phases, there is a need to consider the robot's appearance, capabilities, and behaviors and how the specific ways they are designed can impact the end-user and/or other stakeholders [3, 4]. For example, in the *design* phase, some considerations include: what is the overall purpose of this robot, and who is this for? Additionally, in the *deployment* phase, researchers need to explore how the robot should react in unforeseen situations in the wild. We posit that this long and multifaceted process can be biased by the different stakeholders' own explicit or implicit ableist assumptions which could result in robot designs that reflect ableist paternalism.

On a higher level, the very *idea* of a robot can be paternalistic if it is based on the technoableist assumption that disabled people need their disability to be "fixed". This technoableist bias is illustrated by the concept of the "disability dongle", proposed by Liz Jackson [24]. A disability dongle is a high-tech tool designed for disabled people that fails to respond to their needs and it is biased by misguided assumptions of non-disabled people about the reality of disability. An example of that would be the stair-climbing wheelchairs as they are impressive but impractical, and do not solve the problem of inaccessible spaces.

So, it is clear that there will be multiple stakeholders who will potentially hold different views on how to benefit the user. Additionally, there will be a long series of complicated decision-making processes during all stages of the robot design lifecycle. As a consequence, conflicts are likely to arise, and a roadmap for navigating those negotiations while avoiding the bias of ableist paternalism would be useful.

We propose a three-level risk analysis of ableist paternalism in assistive robotics (Fig. 1). This framing is aimed to convey the multifaceted factors at play in a way that makes it possible to distinguish and evaluate them in their specificity.

In the first level, we propose to look at the *structural factors* that shape decision-making in robot design and implementation. Stakeholders have different interests, power, and knowledge depending on their situation, and the resolution of conflicts and trade-offs that will arise will be affected by these differences. As such, differences in power and the existence of structural inequalities between stakeholders should be explicitly acknowledged [44].




Figure 1: We propose a three-level risk analysis of ableist paternalism in assistive robotics. Our analysis explores the following: (1) Ableism and power inequalities between stakeholders involved in the robot design lifecycle (2) The multi-stage robot design lifecycle (design, development, deployment, and exit), each involving multiple stakeholders (e.g. disabled end-users, robot designers, clinicians, funding agencies) (3) The interactions between robots and disabled end-users which may be a source of conflict and paternalism towards users.

In the second level, we address *design processes*. Currently, there are design methodologies that aim to include users in the decision-making [3] and to align technology-making with social justice concerns [11]. However, the inherent complexity of decision-making in robot design calls for a more detailed investigation of the elements that may give rise to ableist paternalism, even in inclusive design processes. We will focus on the process that results in deciding what is the benefit for the user that the robot is contributing to and the behavior the robot will display to pursue this overall objective. During this process, ableist paternalism may arise depending on how negotiations are navigated and how trade-offs between the interests of stakeholders (including the possibilities of the state-of-the-art technology at hand) are resolved.

In the third level, we deal with the paternalism that can arise in a specific *interaction* between robot and user, that is, in which the robot is performing the paternalistic action itself. Here, we investigate the different ways in which a robot can directly interfere with the user's actions and their implications.

This three-level framing will allow us to point out the different aspects that play a role in potentially paternalistic decisions and actions. We want to be especially mindful of the role of unconscious ableist prejudices, power relationships between stakeholders, and the possibility of conflict between them. Our multilevel analysis aims to clarify the picture, but levels are co-constitutive: structural factors shape design practices, which in turn prompt robot behavior in the interaction with the user.

Given the complexity already mentioned and the sensitive, contextual nature of many health-related decisions, we do not aim to propose a guideline with one-size-fits-all recommendations. On the contrary, we plan to create a checklist for designers. This checklist will be a guiding tool that helps to point out the relevant questions one needs to ask oneself in order to acknowledge and handle the trade-offs related to paternalism. For each level, we will deduce

what are the pertinent questions that would help navigate decision-making. As a result, the checklist would help the HRI community avoid ableist paternalism by making explicit the reasons for each decision, the relevant aspects one should be taking into account and the conflicts that are being negotiated.

4 CONCLUSION

In conclusion, we assert that it is crucial for the HRI community to evaluate the multiple factors that shape robot-mediated paternalism, particularly towards the disabled community. Due to the variety of stakeholders involved, it can be challenging to manage the conflicting interests that may arise from ableist prejudices and asymmetrical power relations.

There has yet to be a clear, in-depth account of how to handle these complex decision-making processes during the robot design lifecycle in relation to disabled people. As part of our ongoing work, we plan to create a usable, practice-oriented guide for robot designers and researchers. It involves a three-level, multifaceted ethical appraisal of robot-mediated paternalism and a checklist that raises questions on how and why certain decisions are being made so as to mitigate ableist paternalism. Ultimately, we hope that the HRI community can draw on our work to understand and counteract the manifold ways that the disabled community's agency and autonomy can be at risk due to robot-mediated ableist paternalism.

REFERENCES

- [1] J. H. Anderson and A. Honneth. Autonomy, vulnerability, recognition, and justice. 2005.
- [2] S. Ashley. *Against Technoableism: Rethinking who Needs Improvement*. W.W. Norton & Company, 2023.
- [3] M. Axelsson, R. Oliveira, M. Racca, and V. Kyriki. Social robot co-design canvases: A participatory design framework. *ACM Transactions on Human-Robot Interaction (THRI)*, 11(1):1–39, 2021.

[4] E. Björling and L. Riek. Designing for exit: How to let robots go. *Proceedings of We Robot*, 2022.

[5] K. R. Bogart and D. S. Dunn. Ableism special issue introduction. *Journal of Social Issues*, 75(3):650–664, 2019.

[6] J. Borenstein and R. Arkin. Robotic nudges: the ethics of engineering a more socially just human being. *Science and engineering ethics*, 22:31–46, 2016.

[7] M. Čačić, J. Avelino, D. Mahr, G. Odekerken-Schröder, and A. Bernardino. Robotic versus human coaches for active aging: An automated social presence perspective. *International Journal of Social Robotics*, 12:867–882, 2020.

[8] F. Cantucci, R. Falcone, and M. Marini. Redefining user expectations: The impact of adjustable social autonomy in human–robot interaction. *Electronics*, 13(1):127, 2023.

[9] A. Chatzidakis, J. Hakim, J. Litter, C. Rottenberg, et al. *The care manifesto: The politics of interdependence*. Verso Books, 2020.

[10] B. Clough. Disability and vulnerability: Challenging the capacity/incapacity binary. *Social Policy and Society*, 16(3):469–481, 2017.

[11] S. Costanza-Chock. *Design justice: Community-led practices to build the worlds we need*. The MIT Press, 2020.

[12] A. Cureton. Some virtues of disability. *International Journal of Applied Philosophy*, 2015.

[13] S. Dodds. Dependence, care, and vulnerability. *Vulnerability: New essays in ethics and feminist philosophy*, pages 181–203, 2014.

[14] G. Dworkin. Defining paternalism. In *New perspectives on paternalism and health care*, pages 17–29. Springer, 2015.

[15] S. Elbeleidy, T. Mott, and T. Williams. Practical, ethical, and overlooked: Teleoperated socially assistive robots in the quest for autonomy. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 577–587. IEEE, 2022.

[16] A. Ferrario, S. Gloclecker, and N. Biller-Andorno. Ai knows best? avoiding the traps of paternalism and other pitfalls of ai-based patient preference prediction. *Journal of Medical Ethics*, 2023.

[17] P. Formosa. Robot autonomy vs. human autonomy: social robots, artificial intelligence (ai), and the nature of autonomy. *Minds and Machines*, 31(4):595–616, 2021.

[18] S. Gillet, M. Vázquez, S. Andrist, I. Leite, and S. Sebo. Interaction-shaping robotics: Robots that influence interactions between other agents. *ACM Transactions on Human-Robot Interaction*, 2024.

[19] A. Hamraie. Crip technoscience. *Building Access: Universal Design and the Politics of Disability*, 2017.

[20] A. Hamraie and K. Fritsch. Crip technoscience manifesto. *Catalyst: Feminism, Theory, Technoscience*, 5(1):1–33, 2019.

[21] W. Hartzog. Unfair and deceptive robots. *Md. L. Rev.*, 74:785, 2014.

[22] T.-H. Ho. Moral difference between humans and robots: paternalism and human-relative reason. *AI & SOCIETY*, 37(4):1533–1543, 2022.

[23] M. Hofmann, D. Kasnitz, J. Mankoff, and C. L. Bennett. Living disability theory: Reflections on access, research, and design. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–13, 2020.

[24] L. Jackson, A. Haagaard, and R. Williams. Disability dongle. *CASTAC Platypus Blog*, 2022.

[25] S. Jayaraman, D. Cruz-Sandoval, A. Kubota, and L. D. Riek. What a health-care worker wants, what a disabled person needs: Exploring stakeholder design tensions in assistive robotics. In *HRI 2024 Workshop on Assistive Applications, Accessibility, and Disability Ethics (A3DE)*, 2024.

[26] A. Kubota, D. Cruz-Sandoval, S. Kim, E. W. Twamley, and L. D. Riek. Cognitively assistive robots at home: Hri design patterns for translational science. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 53–62. IEEE, 2022.

[27] A. Kubota, R. Pei, E. Sun, D. Cruz-Sandoval, S. Kim, and L. D. Riek. Get smart: Collaborative goal setting with cognitively assistive robots. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pages 44–53, 2023.

[28] M. Kühler. Exploring the phenomenon and ethical issues of ai paternalism in health apps. *Bioethics*, 36(2):194–200, 2022.

[29] D. D. Luxton. Ai decision-support: a dystopian future of machine paternalism? *Journal of Medical Ethics*, 48(4):232–233, 2022.

[30] G. Y. K. Ma and W. W. Mak. Maintenance of ableist society for wheelchair users: roles of medical model of disability and witnessed negotiation for accessibility. *Disability & Society*, pages 1–26, 2023.

[31] C. Mackenzie and N. Stoljar. Introduction: autonomy refigured. In *Relational autonomy: Feminist perspectives on autonomy, agency, and the social self*, pages 3–31. Oxford University Press, 2000.

[32] J. Mankoff, G. R. Hayes, and D. Kasnitz. Disability studies as a source of critical inquiry for the field of assistive technology. In *Proceedings of the 12th international ACM SIGACCESS conference on Computers and accessibility*, pages 3–10, 2010.

[33] R. D. Milian and A. Bhattacharyya. Artificial intelligence paternalism. *Journal of Medical Ethics*, 2023.

[34] S. Moharana, A. E. Panduro, H. R. Lee, and L. D. Riek. Robots for joy, robots for sorrow: community based robot design for dementia caregivers. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 458–467. IEEE, 2019.

[35] A. Nanavati, V. Ranganeni, and M. Cakmak. Physically assistive robots: A systematic review of mobile and manipulator robots that physically assist people with disabilities. *Annual Review of Control, Robotics, and Autonomous Systems*, 7, 2023.

[36] M. R. Nario-Redmond, A. A. Kemerling, and A. Silverman. Hostile, benevolent, and ambivalent ableism: Contemporary manifestations. *Journal of Social Issues*, 75(3):726–756, 2019.

[37] L. L. Piepzna-Samarasinha. *Care work: Dreaming disability justice*. Arsenal Pulp Press Vancouver, 2018.

[38] L. D. Riek. Healthcare robotics. *Communications of the ACM*, 60(11):68–78, 2017.

[39] J. L. Scully. Disability and vulnerability: On bodies, dependence, and power. *Vulnerability: New essays in ethics and feminist philosophy*, pages 204–21, 2014.

[40] A. Sharkey and N. Sharkey. We need to talk about deception in social robotics! *Ethics and Information Technology*, 23:309–316, 2021.

[41] T. Titchkosky, E. Cagulada, M. DeWelles, and E. Gold. *DisAppearing: Encounters in Disability Studies*. Canadian Scholars, 2022.

[42] J. C. Tronto. *Caring democracy: Markets, equality, and justice*. NYU Press, 2013.

[43] H. M. Van der Loos, D. J. Reinkensmeyer, and E. Guglielmelli. Rehabilitation and health care robotics. *Springer handbook of robotics*, pages 1685–1728, 2016.

[44] K. Winkle, D. McMillan, M. Arnelid, K. Harrison, M. Balaam, E. Johnson, and I. Leite. Feminist human–robot interaction: Disentangling power, principles and practice for better, more ethical hri. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pages 72–82, 2023.

[45] D. Woody, R. H. Brown, M. Marin, S. Threadcraft, C. P. Harris, J. Syedullah, and M. Ticktin. The politics of care. *Contemporary Political Theory*, 20(4):890–925, 2021.