

# Towards Cognitive Robots That People Accept in Their Home

Nina Moorman, Erin Hedlund-Botti, Matthew C. Gombolay

## Abstract

It is intractable for assistive robotics to have all functionalities pre-programmed prior to deployment. Rather, it is more realistic for robots to perform supplemental, on-site learning about the user's needs and preferences and particularities of the environment. This additional learning is especially helpful for care robots that assist with individualized caregiver activities in residential or assisted living facilities. Each patient is unique, has differing needs, and those needs will change over time, so robots require the ability to adapt and learn. Many assistive robots, ranging in complexity from Roomba to Pepper, have the ability to conduct some of their learning in the home. In this work, we propose to investigate the impacts on end-users of observing this in situ learning through a series of human-subjects experiments. We will assess end-user attitudes towards embodied robots that conduct some learning in the home as compared to a baseline condition where the robot is delivered fully capable. We will additionally compare different modes of learning and interaction to determine whether some are more likely to instill trust.

## Introduction

Care robots currently perform a wide variety of tasks using artificial intelligence (AI) methods that require some degree of learning-at-home. Some of these tasks require observation-based learning; for instance, robotic agents, such as Roomba (iRobot 2002) and Moxi (Robotics 2017) learn the layout of their environment through observation and exploration. Other robots observe users to classify and track user behavior. Robotic agents, such as PHAROS, use this functionality to help monitor user well-being and daily exercise (Martinez-Martin, Costa, and Cazorla 2019). Finally, robots such as the PARO robot perform interaction-based learning of user preferences via physical feedback (Robotics 1998). These forms of supplemental learning enable the robot to adapt to its environment and user(s) and offer care-givers and care-receivers the option of being directly involved in the robot's learning to specify preferences or teach the robot new tasks.

The customization afforded by AI that learns, i.e. using machine learning (ML), is important for ensuring users receive effective, individualized care. Though in situ ML is already in use for some care robots, we lack an understanding

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

of how observing this learning affects user trust. This work aims to develop a better understanding of how users will respond to embodied, learning robots in the home. To do so, we propose to conduct a series of human-subjects experiments that evaluate user trust in four different robot learning methods. These include a baseline, "download" condition, where the robot downloads tasks from the cloud, two forms of learning from demonstration (LfD), and a reinforcement learning (RL) condition. We choose to evaluate these four conditions of learning in order to obtain results that are applicable to the two broad ways in which robotic agents interact with humans while they learn, namely low involvement (i.e., download and RL) and high involvement (i.e., LfD via task demonstration and TAMER (Knox and Stone 2009)). Through surveys and behavioral metrics, we will be able to determine the effect of robot learning on user trust, usability, and adoptability.

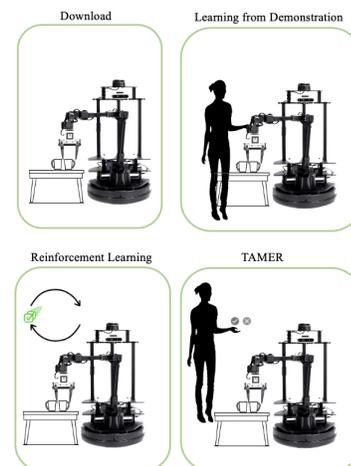


Figure 1: This figure shows the four different robot types.

Informed by this work, we propose to determine the techniques that are most effective in repairing trust in the learning robot. As learning agents may not learn quickly and may fail often, it is important to evaluate the best method of trust repair for embodied, learning agents (Natarajan and Gombolay 2020). We further aim to develop guidelines to inform the design of artificial intelligence for care robotic systems

that operate in residential or nursing-home environments. In our work we propose the following:

1. Study the difference between user trust in a fully pre-engineered robot compared to user trust in a robot that learns in the home.
2. Investigate how to best perform trust repair with respect to embodied robots that learn in the home.
3. Develop guidelines to inform the design of assistive robotic systems deployed in residential environments.

## Related Works

In this section, we describe the prior literature on care robots, acceptability, and trust.

**Care Robots** - Care robots are defined by their function to support care-givers and/or care-receivers (van Wynsberghe 2013). These robots can operate in residential environments where they perform a variety of assistive tasks and promote extended independent living (Johnson et al. 2014; Šabanović et al. 2015; Fiorini et al. 2021; Lee and Riek 2018).

Care robot roles generally fall under physical assistance or medical assistance. Physical assistance includes tasks, such as navigation, fall-prevention, object manipulation, and household chores. Some examples of care robots that perform physical assistance tasks include Moxi (Robotics 2017), Hobbit (Fischinger et al. 2016), Relay (Healthcare 2013), Care-O-Bot 4 (Kittmann et al. 2015), RAMCIP (Kostavelis et al. 2016), and Lio (Mišeikis et al. 2020). Medical assistance includes tasks such as health monitoring, medicine delivery, and the exertion of a social presence for coaching or social interaction (Vitanza et al. 2019; Coradeschi et al. 2013; Umbrico et al. 2020). Examples of care robots that perform medical assistance include Nao (Robotics 2008), Pepper (Robotics 2014), and PHAROS (Martinez-Martin, Costa, and Cazorla 2019).

Whether the care robot performs physical or medical assistance, learning in the home affords the robot an opportunity to observe and adapt to individual user needs and preferences. We seek to understand how this learning impacts users acceptance of these robots.

**Acceptability and Trust** - Acceptability of care robots in depends not only on the benefits the robot can bestow upon the user, but also on the user's perception of and attitudes towards the robot (Cesta et al. 2007). One of the most important attitudes with respect to acceptability is user trust (Yagoda and Gillan 2012; Langer et al. 2019). Trust is defined as a user's attitude that the agent will help them achieve a goal, specifically in a situation of uncertainty, or vulnerability (Kohn et al. 2021; Ullman and Malle 2018).

Prior work in human-automation (HA) trust has categorized trust based on the extent of interaction with the user into dispositional, situational, and learned trust (Hoff and Bashir 2015). These represent baseline trust in automation (Merritt et al. 2013), trust with respect to a particular interaction (Jian, Bisantz, and Drury 2000), and trust developed through a series of interactions (de Visser et al. 2020), respectively. These measures of HA trust can be useful in measuring human-robot (HR) trust, but automation does not always

require a robot embodiment. In our work, we conduct both in-person and remote studies to account for the impact of robot embodiment on trust.

Note that in human robot interactions, anthropomorphism (i.e., the degree to which a robotic agent demonstrates human-like characteristics) has been shown to affect trust (Natarajan and Gombolay 2020). Thus, we will hold the robot's embodiment constant throughout the study and subjectively measure subjects' perception of the robot's anthropomorphism to control for this effect.

In human-robot interaction, there are various mechanisms for studying trust. One approach is to study trust dependent on robot-specific factors, such as performance. This performance-based trust encompasses robot performance and user's awareness of the robot's abilities. In our study, we choose to keep performance consistent between conditions to isolate the impact of observing the learning process. However, there may still be differences in perceived robot ability due to the observation of different forms of learning; as such, we will evaluate performance-based trust. A second lens through which to examine trust is based on interaction-specific factors (i.e., relation-based) (Law and Scheutz 2021). Relation-based trust looks at social factors between the robot and society including robot appearance, adherence to social norms, and morality. We will employ recent work to independently measure these types of HR trust (Malle and Ullman 2021).

## Methodology

In our experiment, participants will watch videos of a robot learning to complete a training task. Then, participants will observe the robot's performance on a related, riskier test task. We will measure their trust in the robot and their perceived risk of the task. This section details the domain, conditions, hypothesis, and metrics of our study design.



Figure 2: This figure shows the experimental set up of the human-subjects study during testing.

## Domain

The domain we choose is depicted in Figure 2. The robot must put a plate together for breakfast, including picking up a knife and cutting a banana in half on a plate, then picking the right medicine and pouring it on the plate. This task is a combination of both a cognitive task/ preparation task (recipe-following) and manipulation task.

While training, the knife used will be a small plastic knife, and the medicine dispensed will be composed of different types of vitamins. During the testing phase, the knife will be a large, sharp, metal knife, and the medicine will be labeled as include morphine, aspirin, and antibiotics. Recall from our definition of trust that we must evaluate trust in a situation of uncertainty, or vulnerability (Kohn et al. 2021; Ullman and Malle 2018). As a result, we will heighten the stakes in the test task to ensure that participants consider their risk tolerance when evaluating the robot.

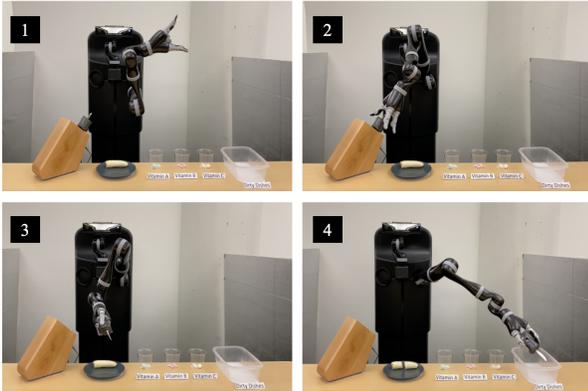


Figure 3: This figure shows a sample training trajectory for the cutting task.

## Conditions

In this subsection, we describe the training video for each of the four conditions. For consistency, the training in all conditions is recorded and viewed in a video format. The way the robot learns depends upon the learning condition.

1. *Download*: In the download condition, the participant observes the robot download the task knowledge from “the cloud.” This serves as the control condition, as no learning is observed.
2. *Reinforcement Learning (RL)*: In the RL condition, the robot shows trial-and-error learning, iteratively learning sub-tasks of the overall task until the task is completed. No explicit reward function is explained, and we intentionally describe stages of the learning vaguely, using terms such as start, middle, and end of training rather than providing iteration count or time metric. We additionally display a pie chart which progressively fills as the training progresses.
3. *Learning from Demonstration (LfD)*: In the LfD condition, the same trial-and-error learning is observed; however, we intersperse videos of a human teaching the robot the sub-tasks, prior to improvement in performance on these sub-tasks at various stages to convey that the robot is acquiring additional knowledge.
4. *TAMER*: Finally, in the TAMER condition, the same trial and error learning is observed; however, we add a video of a human teaching the robot through positive or negative feedback (i.e., binary feedback). This feedback will

be shown to the user in the video through the addition of a graphic of a remote control green and red buttons pressed during training to convey positive and negative feedback, respectively.

## Experiment Populations

Virtual (e.g., via Amazon Mechanical Turk) studies may not fully represent the impact of the robot’s embodiment on trust. However, conducting our study in person would impose a higher health risk and a potentially impractical transportation burden on our target population. As such, we propose to conduct three human-subjects experiments.

*Study 1*: We invite members of the accessible to us at a metropolitan university campus to take our study virtually, on a computer in the laboratory.

*Study 2*: We again invite members of the local population to take our study in the laboratory. In Study 2, the test phase is conducted live, in person, on the Movo Beta robot.

*Study 3*: We invite our target population of caregivers and eldercare nurses in our state to take our study virtually.

We will compare the results of Study 1 and Study 2 to determine the influence of embodied robots on user trust. We will compare the results of Study 1 and Study 3 to determine the difference in trust between caregivers and the general population.

## Hypotheses

**Hypothesis 1** *We hypothesize that participants will trust and adopt learning robots (RL, LfD, TAMER) more than the control robot (download condition). We postulate that participants will understand the learning of a robot that learns visibly over a robot whose learning is not demonstrated.*

**Hypothesis 2** *We hypothesize that participants will trust and adopt LfD more than the other conditions. LfD has been shown to be the most intuitive method of interacting with robotic agents, thus, we hypothesize that participants will demonstrate a preference for this agent type (Fischer et al. 2016; Akgun and Subramanian 2011).*

**Hypothesis 3** *We hypothesize that participants will trust the robots less in Study 2 (embodied robot) as compared to Study 1 (remote). Because participants in Study 2 experience the risk of embodied learning in person but, in Study 1, the task’s risk is virtual, we posit that participants will trust the robot less in person.*

**Hypothesis 4** *We hypothesize that participants will trust the robots less in Study 3 (virtual, with caregivers) as compared to Study 2 (virtual, with local population). We postulate that our population of care providers may find the risk of error to be more tangible and severe.*

## Metrics

We will evaluate a user’s dispositional trust, situational trust and performance-based trust through surveys (Merritt et al. 2013; Jian, Bisantz, and Drury 2000; Malle and Ullman 2021). We will not consider relation-based trust as we will not alter how the robot adheres to social norms. We will also study user trust through behavioral measures (i.e., in terms of reliance on and compliance with the robot) through user

average intervention rates while observing the robot's behavior on the test task of each domain.

1. *Pre-Study Questionnaire* - In the pre-study questionnaire, we will collect demographic information including participant education (Raub 1981), computer science and robotics prior experience (Raub 1981), personality (Donnellan et al. 2006), field of occupation (ACT 2015), and dispositional trust (Merritt et al. 2013). Having observed the unboxing video, the participant additionally provides information regarding their perception of the robot's anthropomorphism (Bartneck et al. 2009) and usability and acceptability (Belanche, Casaló, and Flavián 2012).
2. *Post-Trial Questionnaire* - After the participant observes the training task video, terminated either through interruption or when the video ends, we will ask participants to rate the degree to which they feel the robot accomplished the task, as well as the degree to which they felt the robot was behaving safely.
3. *Post-Study Questionnaire* - In the post-survey questionnaire, we will collect participants' perceived situational trust (Jian, Bisantz, and Drury 2000), performance-based trust (Malle and Ullman 2021), and risk level (Fischhoff et al. 1978). We will also ask two, ad hoc questions. First, we will ask the participant what tasks – from a list of hand-crafted tasks both in and outside of the distribution of tasks observed in the study – they would trust the robot to do. This question measures the extent of user adoption. Secondly, we will ask an open-ended question about the participant's understanding of the robot's learning and perception of robot competence. This second question is to collect qualitative information about user assumptions regarding the robot's learning and competence.

## Procedure

The general procedure of our human-subjects experiments will be as follows. For Study 1 and 2, participants will come in person to the laboratory. Prior to the start of the study, participants will read and sign the consent form. Then, participants will be assigned a unique user ID. Next, participants will watch the unboxing video in which the robotic agent introduces itself and demonstrates its range of mobility and degrees of freedom. After watching the unboxing video, participants will fill out the pre-study questionnaire.

Then, participants will go through the training phase where they will observe the robot learn to perform the task. Ours will be a between-subjects experiment where participants are broken into four conditions (i.e., download, RL, LfD, and TAMER). In the training phase, participants will watch their condition's unique training video. After this training phase, all participants will observe the same final performance video. Note that the only difference between conditions would be the type of robot learning observed: Final performance will be held constant. During the training videos, the cutting task will be performed using a plastic knife and the medications will be different types of vitamins. Next, participants will go through the testing phase. In the testing phase, participants will observe the testing trials where the robot states its goal and demonstrates a suc-

cess rate of 80%. The testing tasks will be performed using a metal knife and the medications will be labeled as antibiotics, aspirin, and morphine. The risk during the testing task is intended to be higher to evaluate how participants trust the robot during high risk situations. During each testing trial, the participant will be instructed to interrupt the robot by clicking a red stop button if they felt that the robot might be acting in an unsafe manner or if they felt that the robot would fail. The interruption data collected here serves to assess reliance. This objective, binary interruption metric, along with the duration of time observing the agent prior to interruption, help to support our findings on trust.

After each testing trial, participants will fill out the post-trial questionnaire where we will ask them to rate the degree to which they trusted the robot to act safely and the degree to which they believe the robot will accomplish the task. The testing iterations will be shown in person in real time in Study 2 and through recorded videos for Study 1 and 3. After the testing phase, the participants will complete the post-study questionnaire.

## Limitations

One limitation of our work is that, in Study 1 and 3, we will be measuring trust based upon a subject's experience watching videos and imagining the robot learning in their home. We aim to quantify the impact of this limitation with Study 2. Another limitation of our work is that we limit our definition of caregiver to nurses employed in assisted living facilities for ease of recruitment in our first investigation. In future work, we plan to increase the breadth of recruitment to include caregivers who are not nurses (e.g., adult children of parents receiving care).

## Future Work

We propose to recruit participants and conduct Studies 1, 2, and 3. Based on these results, we propose a new study (i.e., Study 4) in which we compare various trust repair techniques, applied to the highest effect learning robot from Studies 1-3. We propose to evaluate the following three forms of trust repair established in prior work (de Visser et al. 2020; Baker et al. 2018; Robinette, Howard, and Wagner 2015; Kim et al. 2013).

1. An apology provided directly after the trust violation.
2. Transparency of robot learning through a high-level narration of what is learned.
3. An explanation of what caused the error, without acknowledging fault, provided after the trust violation.

## Conclusion

In this work, we will conduct a series of human-subjects experiments to assess caregiver attitudes towards the concept of embodied care robots that learn in the home, as compared to robots that are delivered fully-capable. Based on the findings of our work, we will develop guidelines that inform the design of care robotic deployed in the home. Finally, we propose to investigate how we can best improve deficient trust in embodied learning robots.

## References

- ACT. 2015. College Majors and Occupational Choices.
- Akgun, B.; and Subramanian, K. 2011. Robot Learning from Demonstration : Kinesthetic Teaching vs . Teleoperation.
- Baker, A. L.; Phillips, E. K.; Ullman, D.; and Keebler, J. R. 2018. Toward an Understanding of Trust Repair in Human-Robot Interaction: Current Research and Future Directions. *ACM Transactions on Interactive Intelligent Systems*, 8(4): 30:1–30:30.
- Bartneck, C.; Kulić, D.; Croft, E.; and Zoghbi, S. 2009. Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots. *International Journal of Social Robotics*, 1(1): 71–81.
- Belanche, D.; Casaló, L. V.; and Flavián, C. 2012. Integrating trust and personal values into the Technology Acceptance Model: The case of e-government services adoption. *Cuadernos de Economía y Dirección de la Empresa*, 15(4): 192–204.
- Cesta, A.; Cortellessa, G.; Giuliani, M. V.; Pecora, F.; Scopelliti, M.; and Tiberio, L. 2007. Psychological Implications of Domestic Assistive Technology for the Elderly. *Psychology Journal*, 5: 229–252.
- Coradeschi, S.; Cesta, A.; Cortellessa, G.; Coraci, L.; Gonzalez, J.; Karlsson, L.; Furfari, F.; Loutfi, A.; Orlandini, A.; Palumbo, F.; Pecora, F.; von Rump, S.; Štimec, A.; Ullberg, J.; and Ötislund, B. 2013. GiraffPlus: Combining social interaction and long term monitoring for promoting independent living. In *2013 6th International Conference on Human System Interactions (HSI)*, 578–585. ISSN: 2158-2254.
- de Visser, E. J.; Peeters, M. M. M.; Jung, M. F.; Kohn, S.; Shaw, T. H.; Pak, R.; and Neerincx, M. A. 2020. Towards a Theory of Longitudinal Trust Calibration in Human–Robot Teams. *International Journal of Social Robotics*, 12(2): 459–478.
- Donnellan, M.; Oswald, F.; Baird, B.; and Lucas, R. 2006. The Mini-IPIP Scales: Tiny-yet-Effective Measures of the Big Five Factors of Personality. *Psychological assessment*, 18: 192–203.
- Fiorini, L.; De Mul, M.; Fabbriotti, I.; Limosani, R.; Vitanza, A.; D’Onofrio, G.; Tsui, M.; Sancarlo, D.; Giuliani, F.; Greco, A.; Guiot, D.; Senegés, E.; and Cavallo, F. 2021. Assistive robots to improve the independent living of older persons: results from a needs study. *Disability and Rehabilitation: Assistive Technology*, 16(1): 92–102. Publisher: Taylor & Francis .eprint: <https://doi.org/10.1080/17483107.2019.1642392>.
- Fischer, K.; Kirstein, F.; Jensen, L. C.; Krüger, N.; Kukliński, K.; aus der Wieschen, M. V.; and Savarimuthu, T. R. 2016. A comparison of types of robot control for programming by Demonstration. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 213–220.
- Fischhoff, B.; Slovic, P.; Lichtenstein, S.; Read, S.; and Combs, B. 1978. How Safe Is Safe Enough? A Psychometric Study of Attitudes Toward Technological Risks and Benefits. *Policy Sciences*, 9: 127–152.
- Fischinger, D.; Einramhof, P.; Papoutsakis, K.; Wohlkinger, W.; Mayer, P.; Panek, P.; Hofmann, S.; Koertner, T.; Weiss, A.; Argyros, A.; and Vincze, M. 2016. Hobbit, a care robot supporting independent living at home: First prototype and lessons learned. *Robotics and Autonomous Systems*, 75: 60–78.
- Healthcare, S. 2013. Autonomous Service Robot For Hospitals - Savioke Relay.
- Hoff, K.; and Bashir, M. 2015. Trust in Automation: Integrating Empirical Evidence on Factors That Influence Trust. iRobot. 2002. iRobot®: Robot Vacuum and Mop.
- Jian, J.-Y.; Bisantz, A.; and Drury, C. 2000. Foundations for an Empirically Determined Scale of Trust in Automated Systems. *International Journal of Cognitive Ergonomics*, 4: 53–71.
- Johnson, D. O.; Cuijpers, R. H.; Juola, J. F.; Torta, E.; Simonov, M.; Frisiello, A.; Bazzani, M.; Yan, W.; Weber, C.; Wermter, S.; Meins, N.; Oberzaucher, J.; Panek, P.; Edelmayer, G.; Mayer, P.; and Beck, C. 2014. Socially Assistive Robots: A Comprehensive Approach to Extending Independent Living. *International Journal of Social Robotics*, 6(2): 195–211.
- Kim, P. H.; Cooper, C. D.; Dirks, K. T.; and Ferrin, D. L. 2013. Repairing trust with individuals vs. groups. *Organizational Behavior and Human Decision Processes*, 120(1): 1–14.
- Kittmann, R.; Fröhlich, T.; Schäfer, J.; Reiser, U.; Weißhardt, F.; and Haug, A. 2015. Let me Introduce Myself: I am Care-O-bot 4, a Gentleman Robot. In *Let me Introduce Myself: I am Care-O-bot 4, a Gentleman Robot*, 223–232. De Gruyter. ISBN 978-3-11-044392-9.
- Knox, W. B.; and Stone, P. 2009. Interactively shaping agents via human reinforcement: The TAMER framework. In *Proceedings of the fifth international conference on Knowledge capture*, 9–16.
- Kohn, S. C.; de Visser, E. J.; Wiese, E.; Lee, Y.-C.; and Shaw, T. H. 2021. Measurement of Trust in Automation: A Narrative Review and Reference Guide. *Frontiers in Psychology*, 12.
- Kostavelis, I.; Giakoumis, D.; Malasiotis, S.; and Tzovaras, D. 2016. RAMCIP: Towards a Robotic Assistant to Support Elderly with Mild Cognitive Impairments at Home. In Serino, S.; Matic, A.; Giakoumis, D.; Lopez, G.; and Ciproso, P., eds., *Pervasive Computing Paradigms for Mental Health*, Communications in Computer and Information Science, 186–195. Cham: Springer International Publishing. ISBN 978-3-319-32270-4.
- Langer, A.; Feingold-Polak, R.; Mueller, O.; Kellmeyer, P.; and Levy-Tzedek, S. 2019. Trust in socially assistive robots: Considerations for use in rehabilitation. *Neuroscience & Biobehavioral Reviews*, 104: 231–239.
- Law, T.; and Scheutz, M. 2021. Chapter 2 - Trust: Recent concepts and evaluations in human-robot interaction. In Nam, C. S.; and Lyons, J. B., eds., *Trust in Human-Robot Interaction*, 27–57. Academic Press. ISBN 978-0-12-819472-0.

- Lee, H. R.; and Riek, L. D. 2018. Reframing Assistive Robots to Promote Successful Aging. *ACM Transactions on Human-Robot Interaction*, 7(1): 11:1–11:23.
- Malle, B. F.; and Ullman, D. 2021. A multidimensional conception and measure of human-robot trust. In *Trust in Human-Robot Interaction*, 3–25. Elsevier. ISBN 978-0-12-819472-0.
- Martinez-Martin, E.; Costa, A.; and Cazorla, M. 2019. PHAROS 2.0—A PHysical Assistant RObot System Improved. *Sensors (Basel, Switzerland)*, 19(20): 4531.
- Merritt, S. M.; Heimbaugh, H.; LaChapell, J.; and Lee, D. 2013. I Trust It, but I Don't Know Why: Effects of Implicit Attitudes Toward Automation on Trust in an Automated System. *Human Factors*, 55(3): 520–534. Publisher: SAGE Publications Inc.
- Mišeikis, J.; Caroni, P.; Duchamp, P.; Gasser, A.; Marko, R.; Mišeikienė, N.; Zwilling, F.; de Castelbajac, C.; Eicher, L.; Früh, M.; and Früh, H. 2020. Lio-A Personal Robot Assistant for Human-Robot Interaction and Care Applications. *IEEE Robotics and Automation Letters*, 5(4): 5339–5346. Conference Name: IEEE Robotics and Automation Letters.
- Natarajan, M.; and Gombolay, M. 2020. Effects of Anthropomorphism and Accountability on Trust in Human Robot Interaction. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, HRI '20*, 33–42. New York, NY, USA: Association for Computing Machinery. ISBN 9781450367462.
- Raub, A. C. 1981. Correlates of Computer Anxiety in College Students.
- Robinette, P.; Howard, A.; and Wagner, A. 2015. *Timing is Key for Robot Trust Repair*.
- Robotics, D. 2017. Moxi. <https://www.diligentrobots.com/moxi>.
- Robotics, P. 1998. PARO Therapeutic Robot. <http://www.parorobots.com/>.
- Robotics, S. 2008. NAO the humanoid and programmable robot. <https://www.softbankrobotics.com/emea/en/nao>.
- Robotics, S. 2014. Pepper the humanoid and programmable robot. <https://www.softbankrobotics.com/emea/en/pepper>.
- Ullman, D.; and Malle, B. F. 2018. What Does it Mean to Trust a Robot? Steps Toward a Multidimensional Measure of Trust. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, HRI '18*, 263–264. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-5615-2.
- Umbrico, A.; Cesta, A.; Cortellessa, G.; and Orlandini, A. 2020. A Holistic Approach to Behavior Adaptation for Socially Assistive Robots. *International Journal of Social Robotics*, 12(3): 617–637.
- van Wynsberghe, A. 2013. Designing robots for care: care centered value-sensitive design. *Science and Engineering Ethics*, 19(2): 407–433.
- Vitanza, A.; D'Onofrio, G.; Ricciardi, F.; Sancarolo, D.; Greco, A.; and Giuliani, F. 2019. Assistive Robots for the Elderly: Innovative Tools to Gather Health Relevant Data. In Consoli, S.; Reforgiato Recupero, D.; and Petković, M., eds., *Data Science for Healthcare: Methodologies and Applications*, 195–215. Cham: Springer International Publishing. ISBN 978-3-030-05249-2.
- Yagoda, R. E.; and Gillan, D. J. 2012. You Want Me to Trust a ROBOT? The Development of a Human–Robot Interaction Trust Scale. *International Journal of Social Robotics*, 4(3): 235–248.
- Šabanović, S.; Chang, W.-L.; Bennett, C. C.; Piatt, J. A.; and Hakken, D. 2015. A Robot of My Own: Participatory Design of Socially Assistive Robots for Independently Living Older Adults Diagnosed with Depression. In Zhou, J.; and Salvendy, G., eds., *Human Aspects of IT for the Aged Population. Design for Aging*, Lecture Notes in Computer Science, 104–114. Cham: Springer International Publishing. ISBN 978-3-319-20892-3.