Not All Errors Are Created Equal: Exploring Human Responses to Robot Errors with Varying Severity

Maia Stiber
mstiber@jhu.edu
Johns Hopkins University
Baltimore, Maryland, USA

Chien-Ming Huang
cmhuang@cs.jhu.edu
Johns Hopkins University
Baltimore, Maryland, USA

ABSTRACT

Robot errors occurring during situated interactions with humans are inevitable and elicit social responses. While prior research has suggested how social signals may indicate errors produced by anthropomorphic robots, most have not explored Programming by Demonstration (*PbD*) scenarios or non-humanoid robots. Additionally, how human social signals may help characterize error severity, which is important to determine appropriate strategies for error mitigation, has been subjected to limited exploration. We report an exploratory study that investigates how people may react to technical errors with varying severity produced by a non-humanoid robotic arm in a *PbD* scenario. Our results indicate that more severe robot errors may prompt faster, more intense human responses and that multimodal responses tend to escalate as the error unfolds. This provides initial evidence suggesting temporal modeling of multimodal social signals may enable early detection and classification of robot errors, thereby minimizing unwanted consequences.

CCS CONCEPTS

• **Human-centered computing**; • **Computer systems organization → Robotics**;

KEYWORDS

Human-robot interaction, social signal, robot error

ACM Reference Format:
Maia Stiber and Chien-Ming Huang. 2020. Not All Errors Are Created Equal: Exploring Human Responses to Robot Errors with Varying Severity. In *ICMI '20: ACM International Conference on Multimodal Interaction, October 25–29, 2020, Utrecht, Netherlands.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/1122445.1122456

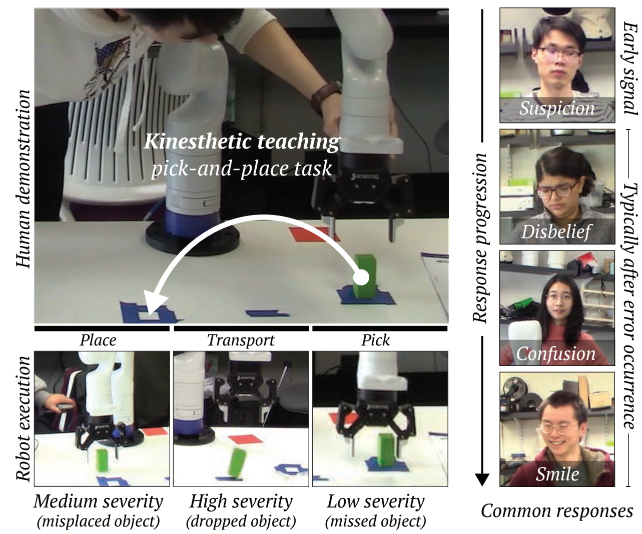

Figure 1: Overview of the experimental conditions and task used in our study, and examples of behavioral responses, exhibited by participants, to unexpected robot errors.

1 INTRODUCTION

While robots promise to positively transform human work and living, they are bound to make errors resulting from imperfect sensing, reasoning, and acting. These robot errors could cause serious physical damage to the surrounding environment and impair people's trust and willingness for adoption. As we continue to develop and integrate robots into human environments, it is important for robots to be able to identify possible errors, allowing for applying subsequent appropriate mitigation strategies to minimize the impact and enabling learning from those errors.

To identify robot errors during situated interactions with people, prior research has investigated how people may respond to a robot's technical and social errors (e.g., [10, 13]) and explored how social signals may be used for error detection [28]. To date, research has mostly used anthropomorphic robots when studying human responses to robot errors; not much has explored how people would exhibit social signals in response to errors produced by non-anthropomorphic robots (e.g., manipulators) that are commonly deployed in workplaces, such as factories. Additionally, how human responses may be used to characterize robot errors, such as severity level, is needed for choosing appropriate error mitigation and has room for exploration. This is especially true when it comes to Programming by Demonstration (*PbD*) scenarios, which have not been examined in this context.

In an effort to fill this knowledge gap, this paper reports an exploratory study of how humans respond to robot errors of varying severity. We contextualize our exploration in a *kinesthetic teaching* setting in which a human provides task demonstrations by directly maneuvering a robotic manipulator (Figure 1). We note that this setting of human teaching robot through physical demonstration is particularly suited to ensure participants have comparable mental

ICMI '20, October 25–29, 2020, Utrecht, Netherlands

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

models of the robot's capability, which is key to studying social responses to robot performance [19]. This setting is also different from prior research that focused on either social interaction scenarios or settings that involve robot guiding human in task completion. In our study, we manipulated robot error severity and measured participants' response time and intensity. Our results indicate that more severe robot errors may prompt faster and more intense human responses. Next, we provide a brief review of relevant prior research that motivates our exploration.

2 BACKGROUND AND RELATED WORK

Errors are inevitable and commonly found in autonomous robot systems [6, 27]. A recent report shows that a robot makes a mistake, negatively impacting task success, every 6 to 20 hours in the field [6]. To better understand and study robot errors, various classifications have been proposed (e.g., [5, 6, 13, 20, 27]). Our exploratory study was focused on technical errors [13, 20].

Though results from prior research agree that robot errors negatively impact task performance and human safety, the effects of robot errors on people's perceptions of the robot—especially perceived trustworthiness—are inconclusive. While some evidence suggested that robots that did not make mistakes were rated significantly more trustworthy than those that did [24], others indicated minor to no statistical significance that errors negatively impacted trust [9, 21]. It was further found that participants liked the robot more when it made mistakes during interactions than when it interacted flawlessly [21], commonly referred as the *pratfall effect*—an increase in likability due to errors [1].

Mixed results were also found in studies that explored the impact of robot error severity on human perceptions. Through a survey and storyboard-based simulation, participants' reactions to low and high error severity were found to be significantly different; moreover, error severity was correlated with loss of trust in the robot [4, 23]. On the other hand, people's trust of robots was found to have no correlation with error severity when humans and robots worked in teams [29]; however, people were more likely to assign blame, usually to themselves or the team as a whole, when experiencing a high severity error than a low one. Error severity, as manipulated by imposing time pressure, was found to influence participants' verbal and gaze behaviors [18]. Our study explored three levels of error severity manipulated by varying task performance and focused on richer social signals, including facial expressions.

Despite the aforementioned mixed effects, errors generally cause negative consequences. Reliable and timely identification of errors is key to successful error mitigation and recovery. Various strategies have been researched to mitigate the impact of robot errors, including asking for help (e.g., [16]) and exhibiting human-like responses such as apologizing (e.g., [12]). If errors can be identified, even after they happen, and the robot employs the appropriate recovery strategy, a positive relationship between people and robots can be formed [31].

A reasonable indicator of errors are social signals [30] as people have been shown to exhibit more behaviors during error situations than error-free ones [5]. In addition, people respond differently when facing social norm violations and technical failures; in particular, technical failures generally resulted in fewer social signals and faster reaction times [20]. Prior works have demonstrated how social signals such as upper body movements (e.g., [28]), gaze (e.g., [2, 3]), and gestures (e.g., [3]) can be used to detect errors effectively. It is worth noting that most of these prior works used human-like robots, which have been shown to elicit different responses to failures than non-humanoid embodiments in social error scenarios [17]. Furthermore, prior research has mostly focused on social interaction scenarios or settings where robots serve as experts or leaders to guide humans through tasks. These settings implicitly relied on participants' existing expectations of robots and thus could possibly influence the participants' social responses. In our study, we explicitly control for people's mental models of the robot by using a kinesthetic teaching scenario in which participants had a practice session to program the robot and formed comparable mental models of the robot's capability.

3 EXPLORATORY STUDY

In this section, we describe an exploratory study that sought to address two research questions: Would people respond socially to robot errors produced by a non-anthropomorphic robot manipulator? How may the social responses, if any, be used to characterize error severity? Answers to these questions will help develop collaborative robots of various forms to interact closely with people.

3.1 Experimental Task and Error Manipulation

The experimental task consisted of waypoint-based kinesthetic demonstration of a pick-and-place task using a Kinova Gen3 robot arm (Figure 1). We designed three conditions corresponding to three levels of error severity: low severity (failing to pick the object), medium severity (placing the object at a wrong location), and high severity (dropping the object during the transportation). The severity manipulation was based on the impact that the error has on the immediate surroundings.

3.2 Study Procedure

Participants were randomly assigned to one of the conditions and asked to program a practice pick-and-place task before the actual experimental task. The practice task was to ensure that the participants knew how to program the robot and that they established comparable mental models of the robot (i.e., the robot was able to execute their program flawlessly). For both the practice and experimental tasks, participants needed to ask the experimenter to run their robot programs. After confirming confidence in programming the robot and programming at least one flawless practice run, the participants were instructed to complete the experimental task. Participants then programmed the experimental task through kinesthetic demonstration using waypoints and signaled to the experimenter to run the program. At this time, rather than running the participant's program, the experimenter executed a program with a pre-programmed error. The entire interaction was recorded to capture how the participants reacted to the robot's erroneous execution of their programs. The participants were fully briefed on the premise of the study and the deception that was involved after the experiment had finished. Participants were then asked to watch their recording and comment on their reaction as part of an open-ended retrospective think-aloud [8]. They also filled out a Big-Five Personality test questionnaire [7] and basic demographics.

3.3 Behavior Analysis and Annotations

To understand social signals exhibited by the participants due to the unexpected errors, we annotated the interaction videos for behavior analysis. A primary coder, who had complete knowledge of the errors and when they occurred during the experiment, performed the coding using VCode [11]. The coder began by identifying the first possible time when the error for each video could be recognized and then annotated behaviors that were a result of the error. In particular, annotated behaviors included verbal (noises or talking), laughing, smiling, looking away from the robot towards the experimenter, scrunching of face, brow moving, and head moving. Based on these annotations, we defined reaction time and reaction intensity, as described below.

Reaction Time (Overt). The reaction time of overt manifestation of the error is the time between when the coder sees the initial overt visual signal of the error and the initial social signal response displayed by the participant. It is the time between the first possible moment the participant could have recognized the error and when the participant reacted to the error. The visual identifier for each type of error differed. The low severity error (robot failing to pick up the object) was signaled by the robot arm gripper moving away from the object without closing the gripper. The medium error (robot placing the object at a wrong location) was indicated by setting the object on the table. The high severity error (robot dropping the object in the middle of the task) was marked by the opening of the gripper in the middle of the robot arm movement.

Reaction Time (Covert). In addition to overt reaction time, we measured covert reaction time, which is the time between when the coder could recognize the error in the program and the first social signal shown by the user. For low and high severity errors, this quantitative metric is the same as the overt version. However, for the medium severity error, which spanned a longer duration, we observed that some participants seemed to be able to detect the error (deviation in trajectory) before the robot placed the object at a wrong location.

Reaction Intensity. We defined intensity by the number of social signals exhibited as a reaction to the error. For example, if the participant reacted to an error with this sequence: brow raising, scrunching, looking towards the experimenter, smiling, issuing a verbal comment, and looking towards the experimenter, then the intensity of the reaction would be six (Figure 2).

3.4 Participants

Seven participants (2 female, 5 male) were recruited for this exploratory study. Of the participants, two were assigned the low, two the medium, and three the high error severity conditions. Participants' ages ranged from 21 to 25 ($M = 23.17, SD = 1.47$). All participants had an engineering background and reported to have considerable experience with robots ($M = 4, SD = 1.41$ in a 5-point Likert scale with 1 being no experience and 5 being a lot of experience) and moderate experience with programming robots ($M = 3.8, SD = 1.30$).

Table 1: Social signals observed in our exploratory study.

Social Signal	Average Duration (*s*) (SD)	Occurrences
Verbal	1.45 (0.96)	9
Look Away	–	7
Smile	5.47 (2.33)	6
Head Movement	1.30 (0.94)	4
Laugh	0.89 (0.84)	3
Scrunch	0.58 (0.07)	3
Brow Raise	2.03 (2.09)	2

3.5 Findings

A total of 34 social signal instances were identified over the seven interaction trials, pertaining to seven different classifications of behavior (Table 1). All errors resulted in social responses from participants with verbal responses, looking away from robot towards experimenter, and smiling being most frequent.

Reaction Times. The average overt reaction time was 0.94 seconds ($SD = 0.62$); covert was 1.31 seconds ($SD = 0.94$). Grouping reaction times by severity, the averages for overt reaction times were 1.31s for low severity ($SD = 1.28$), 1.04s for medium ($SD = 0.14$), and 0.63s for high ($SD = 0.9$), indicating a trend consistent with higher error severity producing faster reactions.

Reaction Intensity. On average, the number of reactions for these errors was 4.86 ($SD = 1.35$). When intensity was broken down by error severity, the average for low severity was 4 ($SD = 1.41$), for medium severity 4.5 ($SD = 2.12$), and for high 6 ($SD = 0$), indicating a trend compatible with more intense reactions being associated with more severe errors.

4 DISCUSSION

In this study, we explored the impact of error severity on multimodal human responses in a human "teach" robot scenario. Our exploration revealed a potential relationship between error severity and reaction time and intensity of response. Below, we discuss our study and results, and their implications for developing collaborative robots capable of detecting errors during situated interactions with people.

Manipulation of Mental Model. Our study involved an explicitly controlled manipulation of the participant's mental model specifically exploring the relationship between error severity and reaction. This setup is different from prior studies where participants' mental models were implicitly assumed appropriate based on the scenario (e.g., [10, 20]). Our explicit, direct manipulation allowed for more pointed results pertaining to errors and human behavioral responses. For example, in a retrospective think-aloud, one participant noted "*for the error part, first of all, I was ... kind of surpris[ed] because based on my first experience [practice task] I think it [the robot] worked perfectly and this robot had built a reputation in my mind. And after the second [actual task], I was like that off, 'like seriously'.*" What the participant was referring to is that the practice task allowed for him to form a mental model that the robot would correctly execute his correct program. However, when the error did occur in the actual task execution, the participant felt surprised and thought that perhaps the robot was inaccurate.

Social Responses to Non-anthropomorphic Robot. As noted before, all participants reacted socially to the unexpected robot

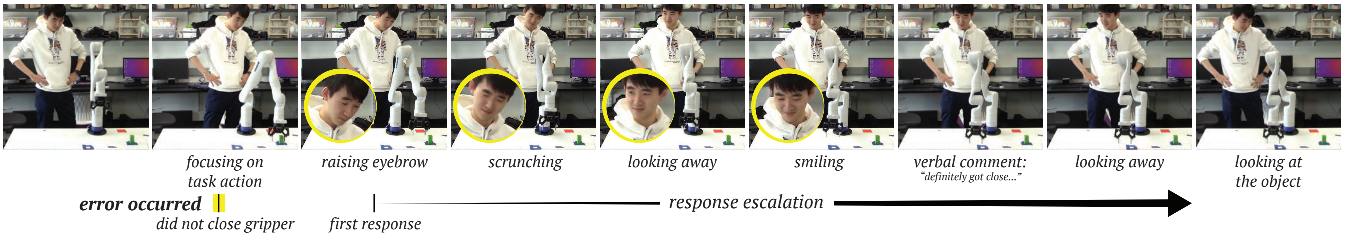

Figure 2: Example of one participant's escalation of behavior after an error of low severity occurred.

errors. This observation contradicts the results from prior studies; specifically, Mirnig *et al.* noted that only 71% of a similar type of errors (right action gone wrong) resulted in social signals [20]. Our observation can possibly be explained by the "Media Equation" theory, which suggests that people react to technologies as they would in human-human social relationships [22]. However, we note that a limitation of this study was that the experimenter was present in the room, which prior work had shown to increase the number of non-verbal signals exhibited [10]. The most frequent social signals observed in our study, namely looking away from the robot towards the experimenter, talking, and smiling, appear to be consistent with prior work [21]. The behavioral response of looking away from the robot towards the experimenter generally happened after robot motion ended and was a sign of participants looking for confirmation and/or error resolution from the experimenter [10].

Escalation of Behavioral Response. In addition to a possible relationship between error severity and reaction time and intensity, we observed how behavioral responses may shift and "escalate" as an error unfolds. As exemplified in Figure 2, a participant's initial reaction tended to be more subtle (i.e., brow movement) and later reaction being more pronounced (i.e., talking in full sentences). While some participants did talk or have short verbal responses (e.g., "*What!*") during the robot movement, they oftentimes "waited" until the end of the robot arm movement to look away from the robot and talk in full sentences. This observation of response progression augments previous findings of frequent combinations of social signal sequences [20]. The sequence of responses and "velocity" of escalation could be modulated by the participants' personalities. Our future work will investigate how temporal modeling of the sequence and escalation of behavioral responses may enable early identification of robot errors.

Expert vs. Novice Participants. Expert and novice participants appeared to have an overall difference in responses. While the group of participants had, as a whole, considerable experience with robots and their programming, their distribution of experience was varied, with two participants having little experience and three having a lot of experience with robots. Notably, three expert users had facial responses to the trajectory of the robot being different than what they had programmed before the error overtly manifested. During one of the expert participant's think-alouds, he commented: "*Did you see what I did with my mouth? I think that is when I noticed that the trajectory is a bit different from what I programmed. But I'm not so sure that happened and I was a bit confused ... Later I confirmed there was a mistake.*" The expert participant was able to identify that there was an error due to trajectory abnormalities that was later confirmed by the actual designated error. This insight indicates that even lower severity errors (e.g., slight deviation of an expected trajectory) than the ones explicitly tested in our study can trigger a visible social signal response. The trajectory deviation did not seem to be noticed by the novice participants.

Personality and Freezing Effect. Personality might shape responses exhibited. While the average reaction time for the study was 0.95 seconds, there was one participant whose reaction time was much longer, at 2.21 seconds. When asked why, the participant said: "*I feel like I recognized the error when the block dropped. It was actually, I'm thinking did I make some mistake or something? So yeah, I was thinking before reacting.*" In other words, while the participant had noticed the error much earlier on, the participant did not react because they were trying to reason about the cause before reacting. This participant is relatively introverted (2.18 in a 5-point Likert scale), based on the Big Five Personality Test. This notion of thinking before reacting may be related to the *freezing* effect reported in previous work where some people stood still after the beginning of an error situation to a behavioral response [10]; the apparent lack of a response was considered a "response to the stress induced by the error situation."

Implications for Developing Collaborative Robots. Findings of our exploratory study illustrate how reaction time and intensity may be leveraged to estimate error severity. Moreover, understanding the sequence of behavioral responses would allow the robot to recognize what stage of reaction the user is in, thereby empowering the robot to tailor its response to the impact of the error as well as user agitation. In addition, encoding information about the human interactant, such as experience with robots and personality, may further enable the robot to adapt error detection sensitivity or size of window to look for reactions to errors.

Future work can extend to additionally include eye gaze in modeling human states, which prior works have suggested to be indicative of robot states (e.g., errors) during situated interactions. Research has shown how gaze cues may be used to understand human intent (e.g., [14]) and robot actions (e.g., [26]), as well as how they can be utilized to enhance robot learning from human demonstration [25], detect errors [2, 3], and facilitate efficient human-robot collaboration [15]. All in all, multimodal human signals present opportunities to enhance human-robot interaction.

ACKNOWLEDGMENTS

This work was partially supported by the Johns Hopkins Institute for Assured Autonomy and the Johns Hopkins University Computer Science Department Fellowship.

REFERENCES

[1] Elliot Aronson, Ben Willerman, and Joanne Floyd. 1966. The effect of a pratfall on increasing interpersonal attractiveness. *Psychonomic Science* 4, 6 (1966), 227–228.

[2] Reuben M. Aronson and Henny Admoni. 2018. Gaze for Error Detection During Human-Robot Shared Manipulation. In *RSS Workshop: Towards a Framework for Joint Action*.

[3] Riccardo Bovo, Nicola Binetti, Duncan P. Brumby, and Simon Julier. 2020. Detecting Errors in Pick and Place Procedures. In *ACM Conference on Intelligent User Interfaces*.

[4] Daniel J. Brooks, Momotaz Begum, and Holly A. Yanco. 2016. Analysis of reactions towards failures and recovery strategies for autonomous robots. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 487–492.

[5] Dito Eka Cahya, Rahul Ramakrishnan, and Manuel Giuliani. 2019. Static and Temporal Differences in Social Signals Between Error-Free and Erroneous Situations in Human-Robot Collaboration. In *Social Robotics*. 189–199.

[6] Jennifer Carlson and Robin R. Murphy. 2005. How UGVs Physically Fail in the Field. *IEEE TRANSACTIONS ON ROBOTICS* 21, 3 (2005), 423–437.

[7] John M. Digman. 1990. Personality Structure: Emergence of the Five-Factor Model. *Annual Review of Psychology* 41 (1990), 417–440.

[8] Sanne Elling, Leo Lentz, and Menno de Jong. 2011. Retrospective Think-Aloud Method: Using Eye Movements as an Extra Cue for Participants' Verbalizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1161–1170.

[9] Rebecca Flook, Anas Shrinah, Luc Wijnen, Kerstin Eder, Chris Melhuish, and Severin Lemaignan. 2019. On the impact of different types of errors on trust in human-robot interaction: Are laboratory-based HRI experiments trustworthy? *Interaction Studies* 20 (2019), 455–486.

[10] Manuel Giuliani, Nicole Mirnig, Gerald Stollnberger, Susanne Stadler, Roland Buchner, and Manfred Tscheligi. 2015. Systematic Analysis of Video Data from Different Human-Robot Interaction Studies: A Categorisation of Social Signals During Error Situations. *Frontiers in Psychology* 6 (2015).

[11] Joey Hagedorn, Joshua Hailpern, and Karrie Karahalios. 2008. VCode and VData: Illustrating a new Framework for Supporting the Video Annotation Workflow. In *Proceedings of the Working Conference on Advanced Visual Interfaces*. 317–321.

[12] Adriana Hamacher, Nadia Bianchi-Berthouze, Anthony G. Pipe, and Kerstin Eder. 2016. Believing in BERT: Using expressive communication to enhance trust and counteract operational error in physical Human-robot interaction. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*.

[13] Shanee Honig and Tal Oron-Gilad. 2018. Understanding and Resolving Failures in Human-Robot Interaction: Literature Review and Model Development. *Frontiers in Psychology* (June 2018).

[14] Chien-Ming Huang, Sean Andrist, Allison Sauppé, and Bilge Mutlu. 2015. Using gaze patterns to predict task intent in collaboration. *Frontiers in psychology* 6 (2015), 1049.

[15] Chien-Ming Huang and Bilge Mutlu. 2016. Anticipatory robot control for efficient human-robot collaboration. In *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*. IEEE, 83–90.

[16] Ross A. Knepper, Stefanie Tellex, Adrian Li, Nicholas Roy, and Daniela Rus. 2015. Recovering from failure by asking for help. *Autonomous Robots* 39 (2015), 347–362.

[17] Dimosthenis Kontogiorgos, Andre Pereira, Boran Sahindal, Sanne van Waveren, and Joakim Gustafson. 2020. Behavioural Responses to Robot Conversational Failures. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, 53–62.

[18] Dimosthenis Kontogiorgos, Sanne van Waveren, Olle Wallberg, Andre Pereira, Iolanda Leite, and Joakim Gustafson. 2020. Embodiment Effects in Interactions with Failing Robots. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery.

[19] Sau-Lai Lee, Ivy Lau, Sara Kiesler, and Chi Yue Chiu. 2005. Human Mental Models of Humanoid Robots. In *IEEE International Conference on Robotics and Automation*. 2767–2772.

[20] Nicole Mirnig, Manuel Giuliani, Gerald Stollnberger, Susanne Stadler, Roland Buchner, and Manfred Tscheligi. 2015. Impact of Robot Actions on Social Signals and Reaction Times in HRI Error Situations. In *International Conference on Social Robotics*. Springer, 461–471.

[21] Nicole Mirnig, Gerald Stollnberger, Markus Miksch, Susanne Stadler, Manuel Giuliani, and Manfred Tscheligi. 2017. To Err Is Robot: How Humans Assess and Act toward an Erroneous Social Robot. *Frontiers in Robotics and AI* 4 (2017).

[22] 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.

[23] Alessandra Rossi, Kerstin Dautenhahn, Kheng Lee Koay, and Michael L. Walters. 2017. How the Timing and Magnitude of Robot Errors Influence Peoples' Trust of Robots in an Emergency Scenario. In *International Conference on Social Robotics*. Springer, 42–52.

[24] Maha Salem, Gabriella Lakatos, Farshid Amirabdollahian, and Kerstin Dautenhahn. 2015. Would You Trust a (Faulty) Robot? Effects of Error, Task Type and Personality on Human-Robot Cooperation and Trust. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. 141–148.

[25] Akanksha Saran, Elaine Schaertl Short, Andrea Thomaz, and Scott Niekum. 2019. Understanding Teacher Gaze Patterns for Robot Learning. (2019). arXiv preprint arXiv:1907.07202.

[26] Alessandra Sciutti, Ambra Bisio, Francesco Nori, Giorgio Metta, Luciano Fadiga, and Giulio Sandini. 2013. Robots can be perceived as goal-oriented agents. *Interaction Studies* 14, 3 (2013), 329–350.

[27] Gerald Steinbauer. 2013. A Survey about Faults of Robots Used in RoboCup. In *RoboCup 2012: Robot Soccer World Cup XVI*. 344–355.

[28] Pauline Trung, Manuel Giuliani, Markus Miksch, Gerald Stollnberger, Susanne Stadler, Nicole Mirnig, and Manfred Tscheligi. 2017. Head and Shoulders: Automatic Error Detection in Human-Robot Interaction. In *ACM International Conference on Multimodal Interaction*.

[29] Sanne van Waveren, Elizabeth Carter, and Iolanda Leite. 2019. Take One For the Team: The Effects of Error Severity in Collaborative Tasks with Social Robots. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, ACM (Ed.). 151–158.

[30] Alessandro Vinciarelli, Maja Pantic, Dirk Heylen, Catherine Pelachaud, Isabella Poggi, Francesca D'Errico, and Marc Schroder. 2012. Bridging the Gap Between Social Animal and Unsocial Machine: A Survey of Social Signal Processing. In *IEEE Transactions on Affective Computing*. 69–87.

[31] Hiroyuki Yasuda and Mitsuharu Matsumoto. 2013. Psychological impact on human when a robot makes mistakes. In *2013 IEEE/SICE International Symposium on System Integration*.