

The Effect of Anthropomorphism and Failure Comprehensibility on Human-Robot Trust

Eileen Roesler¹ & Linda Onnasch²
¹Technische Universität Berlin, Germany
²Humboldt-Universität zu Berlin, Germany

The application of anthropomorphic features to robots is generally considered to be beneficial for human-robot interaction. Although previous research has mainly focused on social robots, the phenomenon gains increasing attention in industrial human-robot interaction, as well. In this study, the impact of anthropomorphic design of a collaborative industrial robot on the dynamics of trust is examined. Participants interacted with a robot, which was either anthropomorphically or technically designed and experienced either a comprehensible or an incomprehensible fault of the robot. Unexpectedly, the robot was perceived as less reliable in the anthropomorphic condition. Additionally, trust increased after faultless experience and decreased after failure experience independently of the type of error. Even though the manipulation of the design did not result in a different perception of the robot's anthropomorphism, it still influenced the formation of trust. The results emphasize that anthropomorphism is no universal remedy to increase trust, but highly context dependent.

INTRODUCTION

The use of robotic systems is becoming increasingly common in various domains, from social assistance to industrial manufacturing. Thereby, robots start across different domains, to interact directly with humans in the same space and time. The fluency and success of this collaboration crucially depends on the human-robot trust (Hancock et al., 2011). Trust in human-technology interaction is most commonly defined as “the attitude that an agent will help achieve an individual's goal in a situation characterized by uncertainty and vulnerability” (Lee & See, 2004, p.51). This attitude is characterized as highly dynamic and varies over time. The phase at the beginning of an interaction can be identified as trust formation (Lewis et al., 2018). Based on factors regarding the robot (e.g. appearance, performance, style of communication), the environment (e.g. task characteristics, team composition, interaction context) and the human (e.g. expertise, cultural background, personality), people develop an initial trust level towards a robot that potentially changes over interaction time and with ongoing experience (Hancock et al., 2011).

A common method to increase the initial trust and support the trust formation in social human-robot interaction (HRI) is the application of anthropomorphic features to the robot, as it is considered to facilitate the growth of meaningful relationships and thereby the acceptance of the robot as a team partner (Duffy, 2003). Anthropomorphism, which can be broadly defined as the human tendency to transfer humanlike characteristics on non-human entities, can be induced not only by a robot's appearance, but also by its movements, style of communication and context (Onnasch & Roesler, in press). While the implementation of anthropomorphic features in social robotics has shown an increase of trust (e.g. Kiesler et al., 2008), there is a lack of research concerning the effects of anthropomorphism in industrial HRI.

This is an important gap, as there is the tendency that industrial robots are no longer separated by fences but increasingly share their workspaces with humans. In contrast to

traditional industrial robots, collaborative robots (Cobots) are intended to interact safely with humans in close proximities. This new form of interaction leads to a change in the perception of the robot from an autonomous tool to a responsive team member. Thus far, the research regarding anthropomorphism in industrial HRI has focused mainly on the transfer of humanlike movements to increase the predictability of robotic actions (e.g. Kuz et al., 2013). Apart from this, the use of anthropomorphic design features in industrial HRI is highly underrepresented. In order to increase trust and acceptance of such a robotic partner, effective design principles of social robotics, like anthropomorphism, might be adopted to industrial HRI.

However, the industrial domain restricts the induction via communication (due to noise level) and often via movements (due to costs of implementation). Nonetheless, the application of anthropomorphism via appearance facets, like moving eyes on a display (e.g. Rethink Robotics Cobot Sawyer) and context modification by framing the robot as team partner with a humanlike name and experience (e.g. Onnasch & Roesler, 2019), seems to be a promising approach to facilitate industrial HRI. More precisely, it is hypothesized that the application of anthropomorphic appearance features and framing (compared to a technical appearance and framing) of an industrial Cobot will lead to higher human trust.

In addition to attribute-based factors of the robot (like anthropomorphism) the human-robot trust development is primarily determined by the performance characteristics of the robot (Hancock et al., 2011). Even though robots are getting more accurate and reliable, they are not operating completely free of errors. This is particularly important regarding the dynamics of trust, as failures might reduce trust in robots considerably (Lewis et al., 2018).

Unfortunately, literature addressing the effects of faulty robots on the dynamics of human trust is still inconclusive. For example, Salem et al. (2015) investigated how erroneous robot behavior in social HRI influenced the perception of the robot and willingness to cooperate with it. The results revealed that participants perceived the robot as less reliable and trustworthy

when it made errors. However, the experience of an error did not affect the willingness to cooperate with the robot. Mainly this finding can be attributed to the artificial nature of the interaction, as the cooperation included highly irrational requests of the robot (e.g. that the subjects should pour orange juice into a plant).

A more realistic experiment in industrial HRI was conducted by Sarkar and colleagues (2017). Thereby participants interacted with an industrial Cobot, which was either functioning perfectly reliable, made one error or made two errors. In contrast to the effects found by Salem and colleagues (2015), they did not find a significant impact of the robot's faults on perceived trustworthiness. Additionally, the number of errors had no influence on the perception of trustworthiness. The authors mainly trace the results to the complexity of the environment and the difficulty of the task.

Surprisingly, the existing results do not clearly support the assumption that the experience of failures leads to a trust dissolution. A reason for this might be that the interactional tasks implemented in the reported studies were either too complex or too unrealistic. In this study, it is therefore aimed to empirically illustrate the general dynamics of human-robot trust from initial trust over trust formation till trust dissolution in a realistic HRI. Moreover, it is assumed that those trust dynamics are dependent on the degree of anthropomorphism induced by robot design. Salem and colleagues (2015) already supposed that "the robot's level of anthropomorphism may lead to different degrees of 'forgiveness' in human interaction partners when errors are displayed" (Salem et al., 2015; p. 319). Unfortunately, it is not possible to relate the different results to the robots' design, as the robots in both studies were equipped with some incomparable anthropomorphic features (e.g. moving digital vs. static physical eyes).

Given the lack of research regarding the relationship of failure experience and anthropomorphic design features on trust, this study investigates how both factors affect the dynamics of trust in industrial HRI. Moreover, the type of failure is further examined by referring to the "easy-error" hypothesis, which states that the comprehensibility of errors made by technology is an important factor that can reduce trust dissolution (Madhavan et al., 2006). It is hypothesized that incomprehensible mistakes made by the robot lead to a stronger decrease in trust compared to comprehensible mistakes.

Additionally, two hypotheses regarding the interactional effect of anthropomorphism and failure comprehensibility are assumed. First, it is hypothesized that incomprehensible mistakes lead to a significantly stronger decrease in trust in the anthropomorphic condition compared to the technical one, as anthropomorphic design modifies the expectations towards the robot. So, for example, participants expect a robot equipped with humanlike eyes to perceive visual input in a humanlike manner (Haring et al., 2013). The implication that an anthropomorphic robot will follow human performance standards further leads to the assumption that errors not detected by the robot will not undermine trust when people expect that these errors might have happened to them as well. Therefore, the second interactional hypothesis states that comprehensible mistakes will lead to a trust decrease in the

technical condition, while they do not influence trust in the anthropomorphic condition.

To investigate the interactional effects of anthropomorphism and failure comprehensibility in a realistic industrial HRI, participants collaborated with a commercially available industrial Cobot in a simple assembly task. The robot either had a technical design and framing or was equipped with anthropomorphic appearance features and presented with an anthropomorphic framing. The dynamics of trust were examined by measuring trust once initially before the actual collaboration started, after a period of perfectly reliable robotic performance and after the experience of failures, which were either comprehensible or incomprehensible.

METHOD

Participants

In total, 50 participants were recruited from Humboldt-Universitaet zu Berlin. Six participants were excluded due to erroneous experimental procedure or because of being outliers on multiple dependent measures. The resulting 44 participants were predominantly students with an age ranging from 18 to 35 years ($M = 24.27$; $SD = 4.6$). The majority identified themselves as female (61.4%) and none of them had previously interacted with the robot used in this study. Participants signed consent forms at the beginning of the experiment and received course credit as compensation at the end of the experiment.

Task and apparatus

The laboratory was arranged as an assembly workspace (Figure 1) with a steel rack containing storage boxes. The industrial Cobot was positioned in front of the rack, facing the human workstation which was set on a high table. The robot used in this experiment was a Sawyer robot from Rethink Robotics (Figure 1) equipped with one arm with 7 degrees of freedom and a range of 1.26 meters.

The main task of the human-robot collaboration was to handover boxes of the steel rack from the robot to the participant. The boxes were first pretentiously scanned by the robot (depicted by a flashlight) to simulate an initial check of the quality (shape, color, size) and quantity of all components inside the box. Afterwards, the robot handed the box to the human coworker. The required movement sequences were programmed by using the software INTERA and included varying movements in the following chronology. First, the robots arm moved over the targeted box and started a flashlight at the gripper to pretend the quality check. Next, the arm moved in a manner that the gripper was able to grab the box. Equipped with the box, the arm moved towards the handover area and waited two seconds in this position before opening the gripper. Subsequently, it moved back to the initial position and waited for fifteen seconds before the start of the next loop. Besides the task relevant movements, no other interaction was possible. The robot's functions and movement patterns were the same in all conditions. The boxes delivered by the robot initiated the task of the participant. The components inside the box were LEGO

bricks to simulate parts of a circuit board, which needed to be assembled in a prescribed style.

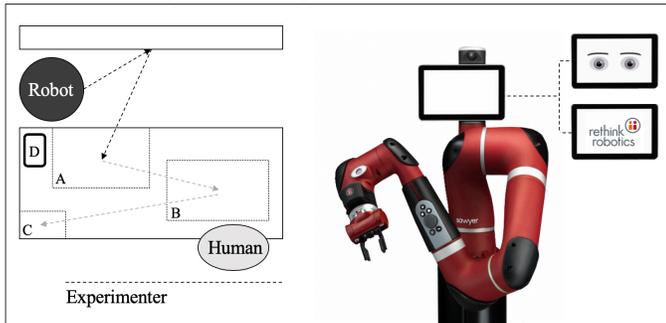


Figure 1. Experimental setting with handover area (A), human workspace (B), submission area (C) and feedback tablet (D) (left) and Sawyer Cobot from Rethink Robotics with both appearance conditions (right).

Afterwards, the participant controlled the quality of the work by comparing the correct location and color of the components to a target circuit board. The visual inspection of the components was implemented as a sequentially redundant task, since the robot checked the color and shape of the components already via the flashlight before transferring the box to the participant. Subsequently, the participant transferred the circuit board to the experimenter, who provided feedback regarding the correctness of the assembled item via a tablet.

Design

The study consisted of a 2 x 2 x 3 mixed design with the two between factors robots design (anthropomorphic vs. technical) and failure comprehensibility (low vs. high) and the within factor experience (initial vs. pre failure vs. post failure).

The different robot design conditions were implemented via appearance (Figure 1) and framing. In the anthropomorphic conditions, the tablet above the robotic arm was equipped with moving eyes and the robot was framed as a colleague named Tom. In contrast, the display showed the Rethink Robotics logo in the technical conditions and the framing characterized the robot as a tool with some technical specifications.

The failures were represented by boxes containing wrong colored LEGO bricks. Since the verification of the components was already done by the robot, the comprehensibility was manipulated by the intensity of the color deviation. In the incomprehensible conditions, one brick massively deviated from the target color (dark grey vs. pink), whereas in the comprehensible conditions one brick marginally differed from the target color (dark grey vs. light grey).

Dependent measures

The key dependent variable was trust in the robot. It was subjectively measured using the Trust in Industrial Human-Robot Collaboration questionnaire (Charalambous et al., 2016). The questionnaire consisted of ten items regarding the perception of motion, speed, safety and gripper reliability, which were rated on a five-point Likert scale and accumulated for the overall score. Besides the specific aspects (mainly focused on safety and technical aspects) of the questionnaire,

two single item questions regarding the overall trust in the robot and the perceived reliability of the robot were used. Both ratings were given in percent.

Behavioral data on how long it took participants to assemble and check the circuit board were derived from video recordings. It was assumed that higher trust leads to a less intense quality check of the components (e.g. color), as this task was already fulfilled by the robot.

To prevent confounding effects of participants' attitudes towards technology and robots in particular, the affinity of technology (Franke et al., 2019), confidence in technology (Feuerberg et al., 2005), tendency to anthropomorphize (Waytz et al., 2010) and negative attitudes towards robots (Syrdal et al., 2009) scales were used as control variables.

To check whether the manipulation of anthropomorphism was successful the Godspeed questionnaire (Bartneck et al., 2009) and the revised Godspeed questionnaire (Ho & MacDorman, 2010) were used. Both questionnaires consist of semantic differential items, which were rated on a five-point Likert scale. The manipulation of failure comprehensibility was checked by asking the participants to rate on a five-point Likert scale whether they too could have committed the failure.

Procedure

All subjects were randomly assigned to one of the four conditions and received corresponding written instructions including the framing of the robot. After filling out sociodemographic questions, the participants completed questionnaires regarding the control variables. Subsequently, the participants were exposed to the robot and the display was either showing eyes or the company's logo, according to the respective condition. The appearance of the static robot was then assessed via anthropomorphism questionnaires. Before the actual task started, the trust questionnaires were filled out (initial measure). Trust was measured again after completing the first eight faultless trials (pre-failure measure). Afterwards, the next eight trials contained either two comprehensible or two incomprehensible failures at trial eleven and trial fourteen. After the last trial (trial 16), trust was measured again (post-failure measure) and the participants were asked, whether they too could have committed the failure. The objective trust measurement was calculated by using video data of the first two trials (initial trust), the two trials before the second subjective trust measurement (pre-failure trust) and the last two trials (post-failure trust). The entire experimental procedure lasted approximately 60 minutes, whereby the interaction with the robot lasted about two thirds of the time.

RESULTS

Attitudes towards technology

First, the variables regarding the individual attitude concerning technology and anthropomorphism were analyzed between the four conditions using a one-way ANOVA. The analyses revealed no significant differences between the four groups in the affinity to technology, $F(3, 40) = 0.56, p = .65$;

confidence in technology, $F(3, 40) = 0.11, p = .95$, tendency to anthropomorphize, $F(3, 40) = 0.06, p = .98$ and negative attitudes towards robots, $F(3, 40) = 0.25, p = .86$.

Manipulation check

Surprisingly, the scores of the anthropomorphism scale (Godspeed questionnaire) revealed comparable ratings in the anthropomorphic ($M = 1.92, SD = 0.64$) and technical ($M = 1.74, SD = 0.48$) condition. This outcome is supported by the nearly identical ratings of the human likeness indices (revised Godspeed questionnaire) of the anthropomorphic ($M = 1.92, SD = 0.55$) and technical ($M = 1.96, SD = 0.66$) condition. Independent-samples t-tests revealed neither significant differences between the technical and anthropomorphic condition on the anthropomorphism score, $t(42) = 1.06, p = .3$, nor on the humanness dimension, $t(42) = -0.25, p = .81$.

The check for the failure comprehensibility revealed a successful manipulation. Participants rated the incomprehensible failure ($M = 2.39, SD = 1.53$) as less understandable than the comprehensible one ($M = 3.62, SD = 1.4$). This difference between failure conditions was statistically significant, as an independent-samples t-test revealed, $t(42) = -2.77, p < 0.01$.

Trust development

Effects on subjective and objective trust were analyzed by a 2 (robot design) x 2 (failure comprehensibility) x 3 (experience) mixed ANOVA with repeated measurements for the last factor.

The analyses of the trust questionnaire revealed a significant main effect of experience, $F(2, 39) = 6.33, p < .01, \eta^2 = .25$. Post hoc tests using Bonferroni correction for multiple comparisons showed that participants' trust was significantly lower ($p < .01$) prior to interaction ($M = 39.89, SE = 0.61$) than after faultless interaction ($M = 42.16, SE = 0.6$).

The analysis of single item trust too revealed a significant main effect of experience, $F(2, 39) = 11.72, p < .01, \eta^2 = .38$. Again, post hoc tests using Bonferroni correction for multiple comparisons revealed that initial trust ($M = 73.93, SE = 2.57$) was significantly lower ($p < .01$) than trust after faultless interaction ($M = 80.16, SE = 2.17$). Additionally, trust after the experience of the robot's fault ($M = 73.3, SE = 2.58$) was significantly lower ($p < .01$) than trust before this experience. Moreover, the main effect of anthropomorphism, $F(1, 40) = 3.65, p = .06, \eta^2 = .08$, just failed to reach the conventional level of significance. Unexpectedly, participants tended to express higher trust in the technical robot condition ($M = 80.05, SE = 3.14$) compared to the anthropomorphic robot condition ($M = 71.56, SE = 3.15$).

The analysis of the single item perceived reliability (Figure 2) showed a significant main effect of experience, $F(1.49, 59.41) = .15.84, p < .01, \eta^2 = .28$. Significant differences existed between the initial rating ($M = 81.16, SE = 1.9$) and the post-failure rating ($M = 70.82, SE = 2.86$) and between the pre-failure ($M = 83.58, SE = 2.03$) and the post-failure rating, as post hoc tests using Bonferroni correction for multiple

comparisons revealed (both $p < .01$). Furthermore, the analysis showed a significant main effect of anthropomorphism, $F(1, 40) = 6.12, p < .05, \eta^2 = .13$ and, again, the ratings were higher in the technical robot condition ($M = 83.06, SE = 2.59$) than in the anthropomorphic robot condition ($M = 73.98, SE = 2.06$).

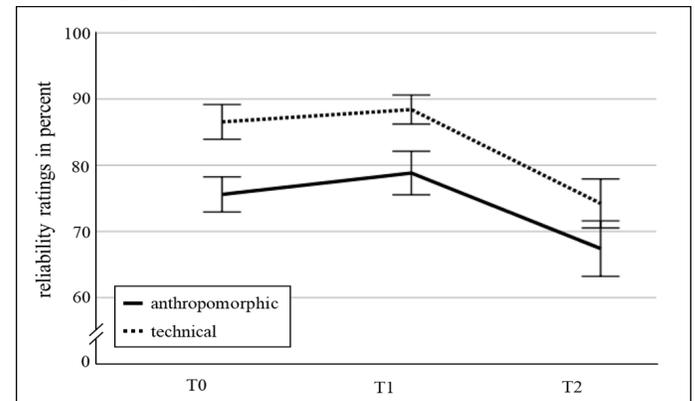


Figure 2. Participant's perceived reliability mean ratings and standard error prior the interaction (T0), after faultless interaction (T1) and faulty interaction (T2) for both design conditions.

The analyses of the assembly time, as objective indicator for trust, showed a significant main effect of experience, $F(1.55, 57.4) = 92.59, p < .01, \eta^2 = .71$. Post hoc tests using Bonferroni correction for multiple comparisons revealed that participants needed significantly longer time to assemble and verify the parts ($p < .01$) in the first interactions ($M = 17.71, SE = 0.58$), compared to the assembling of the box after faultless ($M = 12.67, SE = 0.4$) and faulty interaction ($M = 12.08, SE = 0.36$). Neither the main effects of robot design and failure type nor any of the interaction effects became significant (all $p > .1$).

DISCUSSION

The objective of the presented study was to examine the joint effects of anthropomorphic robot design and the experience of more or less comprehensible failures on human trust in a realistic industrial human-robot collaboration.

Based on previous assumptions in social HRI (Duffy, 2003), it was supposed that the trust in the robot would be higher if the robot was designed in a more anthropomorphic manner. Surprisingly, the results did not confirm this assumption. In fact, participants tended to show higher trust ratings in the technical robot conditions compared to the anthropomorphic robot conditions. This trend of negative effects of anthropomorphic design features on human-robot trust was further supported by the reliability ratings, as participants perceived the robot as significantly less reliable if it was designed in a more anthropomorphic way. This is particularly remarkable, as participants did not distinguish between the technical and anthropomorphic conditions in their perceived degree of human likeness. The latter result might be due to the fact that the introduction of only eyes as an anthropomorphic appearance feature in combination with a framing was not sufficient to impact the overall rather technical appearance of the robot to a degree differentiated by the questionnaires. However, in line with previous findings (Onnasch & Roesler, 2019), even if the manipulation of

anthropomorphism of the physically same robot did not modify the perception of the human likeness, it still influenced the attitudes towards the robot. Specifically, the anthropomorphic design features might have masked the relevance of the robot as a precise working tool and thereby undermined the perceived reliability of the robot. The effect might also be related to the cartoon-like design of the eyes, which could be perceived as inappropriate in professional scenarios.

Furthermore, it was assumed that incomprehensible failures would lead to a more pronounced decline in trust than comprehensible ones. However, the results provide no such effect for either of the dependent variables. This is especially surprising, as participants rated the obvious failure as significantly less understandable. However, the effect might have been covered by the overall high effect of trust dissolution after an initial failure experience, as the failure occurrence lead to a pronounced decrease in perceived reliability and trust.

Contrary to the expectations, the operationalization of objective trust showed a significantly longer assembly time in the first two trials compared to the other times of measurement, but no effect of the design of the robot or failure experience. This can be attributed mainly to a training effect of the task.

CONCLUSION

Altogether, the present study provides initial insights into the consequences of non-task related anthropomorphic design in industrial HRI. In contrast to the beneficial use of predictive and therefore task-related anthropomorphic cues (e.g. Kuz et al., 2013), the results suggest a negative effect of the usage of non-predictive anthropomorphism. The findings of this study emphasize that anthropomorphic design should not be generally applied to industrial robots if it does not support the coordination of the interaction partners. Future research should systematically examine the effect of various anthropomorphic design features by taking the context sensitivity into consideration.

In addition, the study empirically illustrates the dynamics of the development of trust - from trust formation to trust dissolution due to failure experience. Based on these findings, future research should go further by taking into account trust restoration

In summary, anthropomorphic design does not facilitate industrial HRI indiscriminately, and instead leads to a reduction of the perceived reliability of the robot. As found by the meta-analyses by Hancock et al. (2011), performance-based factors mostly determine the development of trust in HRI. The application of anthropomorphic design independently from the actual performance of the robot, can undermine the perceived reliability and therefore hinder the adequate formation of trust. Beyond this, the results do not indicate that anthropomorphism makes people more forgiving towards robots' failures.

REFERENCES

Bartneck, C., Kulić, D., Croft, E., & 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. <https://doi.org/10.1007/s12369-008-0001-3>

- Charalambous, G., Fletcher, S., & Webb, P. (2016). The development of a scale to evaluate trust in industrial human-robot collaboration. *International Journal of Social Robotics*, 8(2), 193–209. <https://doi.org/10.1007/s12369-015-0333-8>
- Duffy, B. R. (2003). Anthropomorphism and the social robot. *Robotics and autonomous systems*, 42(3-4), 177-190. [https://doi.org/10.1016/S0921-8890\(02\)00374-3](https://doi.org/10.1016/S0921-8890(02)00374-3)
- Feuerberg, B.V., Bahner, J.E. & Manzey, D. (2005). Interindividuelle Unterschiede im Umgang mit Automation - Entwicklung eines Fragebogens zur Erfassung des Complacency. In L. Urbas & C. Steffens (Eds.), *Zustandserkennung und Systemgestaltung*. 6. Berliner Werkstatt Mensch-Maschine-Systeme (pp. 199-202). VDI-Verlag.
- Franke, T., Attig, C., & Wessel, D. (2019). A personal resource for technology interaction: development and validation of the affinity for technology interaction (ATI) scale. *International Journal of Human-Computer Interaction*, 35(6), 456-467. <https://doi.org/10.1080/10447318.2018.1456150>
- Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., De Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors*, 53(5), 517-527. <https://doi.org/10.1177/0018720811417254>
- Haring, K. S., Watanabe, K., & Mougnot, C. (2013). The influence of robot appearance on assessment. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 131–132). IEEE. <https://doi.org/10.1109/HRI.2013.6483536>
- Ho, C. C., & MacDorman, K. F. (2010). Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices. *Computers in Human Behavior*, 26(6), 1508–1518. <https://doi.org/10.1016/j.chb.2010.05.015>
- Kiesler, S., Powers, A., Fussell, S. R., & Torrey, C. (2008). Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition*, 26(2), 169–181. <https://doi.org/10.1521/soco.2008.26.2.169>
- Kuz, S., Mayer, M. P., Müller, S., & Schlick, C. M. (2013). Using anthropomorphism to improve the human-machine interaction in industrial environments (Part I). In *International Conference on Digital Human Modeling and Applications in Health, Safety, Ergonomics and Risk Management* (pp. 76-85). Springer. https://doi.org/10.1007/978-3-642-39182-8_9
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- Lewis, M., Sycara, K., & Walker, P. (2018). The role of trust in human-robot interaction. In *Foundations of trusted autonomy* (pp. 135-159). Springer. https://doi.org/10.1007/978-3-319-64816-3_8
- Madhavan, P., Wiegmann, D. A., & Lacson, F. C. (2006). Automation failures on tasks easily performed by operators undermine trust in automated aids. *Human Factors*, 48(2), 241–256. <https://doi.org/10.1518/00187200677724408>
- Onnasch, L., & Roesler, E. (2019). Anthropomorphizing Robots: The Effect of Framing in Human-Robot Collaboration. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 63(1), 1311–1315. <https://doi.org/10.1177/1071181319631209>
- Onnasch, L., & Roesler, E. (in press). A taxonomy to structure and analyze human-robot interaction. *International Journal of Social Robotics*.
- Salem, M., Lakatos, G., Amirabdollahian, F., & Dautenhahn, K. (2015). Would you trust a (faulty) robot? Effects of error, task type and personality on human-robot cooperation and trust. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 1-8). IEEE. <https://doi.org/10.1145/2696454.2696497>
- Sarkar, S., Araiza-Illan, D., & Eder, K. (2017). Effects of faults, experience, and personality on trust in a robot co-worker. arXiv preprint arXiv:1703.02335.
- Syrdal, D. S., Dautenhahn, K., Koay, K. L., & Walters, M. L. (2009). The negative attitudes towards robots scale and reactions to robot behaviour in a live human-robot interaction study. Adaptive and Emergent Behaviour and Complex Systems - *Proceedings of the 23rd Convention of the Society for the Study of Artificial Intelligence and Simulation of Behaviour*, AISB 2009, (pp. 109–115). SSAISB.
- Waytz, A., Cacioppo, J., & Epley, N. (2010). Who sees human? The stability and importance of individual differences in anthropomorphism. *Perspectives on Psychological Science*, 5(3), 219–232. <https://doi.org/10.1177/1745691610369336>