The Effect of Robot Decision Making on Human Perception of a Robot in a Collaborative Task - A Remote Study

Ali Noormohammadi, Abhinav Dahiya, Alexander M. Aroyo, Stephen L. Smith, Kerstin Dautenhahn
University of Waterloo, Faculty of Engineering
Canada

ABSTRACT

The use of collaborative robots is becoming more widespread across industries. This makes it essential to study robot planning in order to work effectively and smoothly with human teammates while maintaining a positive human perception of the robots. This paper evaluates the influence of a robot's strategy and decision making on the participants' perception of the robot. We designed an online experiment where a robot and participants need to collaborate and organize a set of objects. We studied three different strategies where the robot either prioritizes the human's objective, its own objective, or uses a balanced strategy. We then analyze and report the results based on participants' answers to questionnaires before and after the experiment, their comments, and their actions during the experiment. The results show that strategies prioritizing the human's objective, or balancing between the robot's and the human's objectives can effectively improve participants' perception of the robot and create a collaborative environment.

CCS CONCEPTS

• **Computing methodologies → Robotic planning**.

KEYWORDS

Collaborative scenario, Robot decision making, Human's perception of the robot

ACM Reference Format:

Ali Noormohammadi, Abhinav Dahiya, Alexander M. Aroyo, Stephen L. Smith, Kerstin Dautenhahn. 2021. The Effect of Robot Decision Making on Human Perception of a Robot in a Collaborative Task - A Remote Study. In *Proceedings of the 9th International Conference on Human-Agent Interaction (HAI '21), November 9–11, 2021, Virtual Event, Japan.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3472307.3484671

1 INTRODUCTION

A quick search of the term "Industrial Robots" returns a plethora of results showing large, fast, and powerful traditional industrial robotic arms. These traditional robotic arms, one of the most successful commercial robots, can work quickly, efficiently, and safely twenty-four hours a day for mass production. However, the Industry 4.0 era directs manufacturers to also employ robots that are able to work closely and collaborate safely with humans, an ability that traditional industrial robots lack [1, 9]. Such human and robot collaboration can compensate for a human's limitations (e.g., repeatability, precision, speed, cognitive load, physical strength, and stamina) while using their superior abilities (e.g., adapting and coping with the uncertainties of customers' changing preferences, or changing environments and situations that require human-level intelligence and decision-making).

In this paper, we investigate how a cobot's strategy influences the human perception of the robot and its collaboration. In the experimental scenario, having a single-human collaborating with a single-robot, the following three different strategies are adopted: the robot either prioritizes objectives of the human teammate, prioritizes own objectives, or uses a balanced strategy (i.e., considers both human and robot objectives). Then, based on the data collected from the questionnaires and participants' actions, we analyze and discuss influences of each strategy on their perception of the collaboration and the robot.

1.1 Related work

Cobot programming: In a scenario where a robot works in isolation or without any direct interaction with humans, the control objectives are based on the robot's performance metrics — time, energy, traversed distance, covered area, etc. However, in HRC, humans may take a collaborative or even a leading role, and the control objectives thus need to be modified to meet the humans' comfort needs, both physical and mental. Many studies have focused on the human teammate's physical comfort, such as their ergonomic posture [5] and fatigue levels [20]. Trust [15], intention [14, 22], knowledge of the task [2, 8], and acceptance of cobots are some key topics of interest in HRC specifically, and human-robot interaction (HRI) in general. Programming robots to consider humans' mental states in their decision making can significantly enhance the quality and efficiency of collaboration and humans' perception of robots [2, 6, 19].

Trade-offs between human and task objectives: Clearly, deploying cobots in industry is economically justifiable only when they can offer long- or short-term benefits and returns on investment [9]. Hence, HRC systems have to not only improve and maintain humans' perception of robots at a reasonably high level, but need to improve task efficiency as well. In [11], a probabilistic approach was adopted to predict human actions in order to optimize a cobot's assitance and reduce wait times. In [12], workloads per sub-task were minimized by distributing assembly sub-tasks between a human and a cobot. In another study [4], an optimal task distribution between humans and robots was proposed to minimize the completion time of a realistic production process.

HAI '21, November 9–11, 2021, Virtual Event, Japan

ACM ISBN 978-1-4503-8620-3/21/11.
https://doi.org/10.1145/3472307.3484671

However, the problem is more challenging with cases in which human objectives and task objectives conflict, and pull in opposite directions, or where human's plans and actions are not optimal. To cope with these problems, the robot needs to balance the two objectives and adopt a policy to guide the human to make (sub)optimal policies and decisions. Probabilistic and game theory approaches are commonly used to achieve a mutual benefit, or to guide (persuade) adaptable agents toward (sub)optimal actions [10, 16, 18].

For these cases, little is known about the influence of robot's strategy on human perception of the robot. Some studies have focused on the effects of robot's blame attribution (blaming either itself, the human, or the human-robot team) in HRC after the collaboration or after doing its actions [3, 13, 21]. The authors in [7] have investigated how a robot's prosocial and selfish behaviors in a collaborative game can affect the human's perception of competence, responsibility, trust, and preference for future collaborations. However, our study also considers a balanced behavior to compromise between the robot and the human's objectives and priorities.

2 RESEARCH QUESTIONS AND HYPOTHESIS

This study seeks to learn how a robot's decision making in a collaborative task can affect its human teammate's perception of the robot. To do so, we consider three different strategies for the robot: (i) **Strategy 1 (SH)** - The robot completely prioritizes the human's task and objective; (ii) **Strategy 2 (SR)** - The robot completely prioritizes its own task and objective; (iii) **Strategy 3 (SHR)** - The robot adopts a compromising strategy that balances its own and the human's priorities.

The first two strategies are at opposite ends of the planning spectrum, where we expect to see the highest and the lowest levels of satisfaction and trust. Based on these three strategies, we developed the following hypothesis:

Hypothesis 1: The human's (a) trust in the robot, (b) perception of the robot's performance, (c) perception of the robot's collaboration, and (d) willingness to collaborate with the robot, will all be the highest when the robot uses strategy 1 (SH), second highest when using strategy 3 (SHR), and lowest when using strategy 2 (SR).

The study received ethics approval from the University of Waterloo Research Ethics Board (ORE #42760).

3 EXPERIMENTAL METHOD

For this experiment, we have designed an online interactive simulation environment in which the human and the robot need to organize a certain number of objects.

Materials and setup: For the simulation environment, we chose Webots robot simulator and the UR5 robot, a well-known 6-degree-of-freedom cobot manufactured by Universal Robots.

Experiment task: Participants collaborate with a robot to organize certain objects. As shown in Figure 1, the human (1) stands beside their table (2). Their task is to put the colored blocks on the corresponding squares marked on the board (3) placed beside them in the shortest time. Participants have to fill each row before moving to the next one, starting with the first one and end at the fourth. The color of the blocks and the squares have to be matched. For example, a blue block has to be placed on a blue square. The robot (5) also needs to put orange objects in the orange bin, next

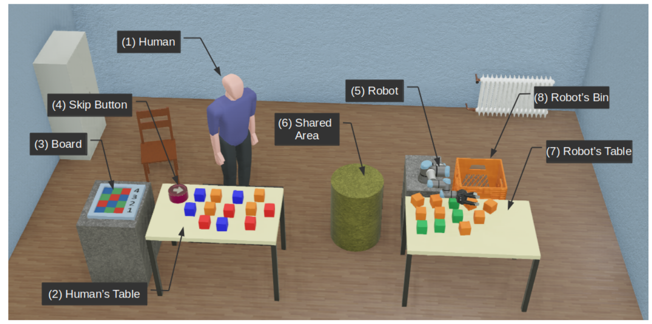

Figure 1: Online experiment environment

to it. This is a turn-taking activity, meaning that the robot takes the first turn, and when it is done, the human chooses their action; then, the robot makes its next move, and so on.

There are four red, four blue, and four orange blocks on the human's table (2). Participants need blue, red, and green blocks to complete the board, but they only have access to blue and red ones on their table. Thus, the robot has to place the green objects on the shared space (6), the green platform beside their table, so that participants can use them to complete the board. Participants also need to place the orange blocks from their table on the same shared area (6) so that the robot can take them and put them in its bin (8). The shared space has a capacity of two objects. There is also a skip button (4) on the participants' table, which they can press if they want to skip their turn. The participants' goal is to fill the board in the shortest time, while the robot's goal is to put all the orange blocks in the bin. The task finishes when all objects are removed from the tables and the shared space. Based on the explained scenario, there is a set of possible actions for both human and robot, from which they can choose a feasible action in each turn:

Human's possible actions A^h: (1) Pick up a blue block from the table and place it on the board. (2) Pick up a red block from the table and place it on the board. (3) Pick up a green block from the shared space and place it on the board. (4) Pick up an orange block from the table and place it on the shared space. (5) Skip their turn.

Robot's possible actions A^r: (1) Pick up an orange block from the table and place it in the bin. (2) Pick up an orange block from the shared space and place it in the bin. (3) Pick up a green block from the table and place it on the shared space. (4) Skip its turn.

System Modeling and Planning: The human-robot team can be represented by the system state $\mathbf{x}_k \in X$, robot action $a^r \in A^r$, and human action $a^h \in A^h$. After each robot action, the robot receives a real-valued reward $r(\mathbf{x}_k, \mathbf{u}_k)$, where $\mathbf{u}_k = \left[a^r, a^h\right]$ represents the vector of the system input, at time step k.

At any time step k, first the robot takes its action a^r_k changing the system state to an intermediate state $\mathbf{x}'_k$. The human observes the robot's action and then decides their action a^h_k, transitioning the system to the next state $\mathbf{x}_{k+1}$. Thereafter, the robot takes its action for time step $k+1$ and so on. This turn-taking between the robot and the human continues until the system reaches the goal state.

In this experiment, the human policy is not known in advance and no online or offline estimation is made during the robot planning. Instead, it is assumed that the human takes the optimal action at every step, i.e, the one that helps maximizing the accumulated reward. Define the optimal policy $\pi^*: X \to U$ that maximizes the discounted sum of rewards as,

$$\pi^*(\mathbf{x}) = \arg\max_{\pi \in \Pi} \left[\sum_{k=0}^{\infty} \gamma^k r\left(\mathbf{x_k}, \pi(\mathbf{x}_k)\right) \,\middle|\, \mathbf{x}_0 = \mathbf{x} \right], \tag{1}$$

where Π is the set of all admissible policies $\pi : X \to U$.

During the experiment, at each time step, the robot chooses its action according to the policy π^*. The reward is only the function of the current state of the system and the robot's action, and is a weighted combination of two scalar functions $\varphi_R(\mathbf{x_k}, a_k^r)$ and $\varphi_H(\mathbf{x_k}, a_k^r)$ reflecting the extent to which the action a_k^r aims to satisfy the robot's and the human's preferences:

$$r\left(\mathbf{x}_k, a_k^r\right) = (1-q) \times \varphi_R(\mathbf{x_k}, a_k^r) + q \times \varphi_H(\mathbf{x_k}, a_k^r), \tag{2}$$

where $q \in [0, 1]$ is a constant value indicating the importance of the human preferences. The three strategies can be implemented by simply setting the value of q. For strategy 1 (pure-assistive) we set $q = 1$, for strategy 2 (self-serving) we set $q = 0$ and for strategy 3 (balanced) we set $q = 0.5$.

Procedure: We recruited 50 participants (48 graduate and 2 undergraduate), through email advertisements. The cohort for each condition (planning strategy) consisted of: (i) Condition-1: 17 participants (9 men and 8 women) with an average age of 28.71 ± 9.59 years (ii) Condition-2: 17 participants (10 men and 7 women) with an average age of 28.07 ± 4.18 years (iii) Condition-3: 16 participants (9 men and 7 women) with an average age of 27.56 ± 4.76 years. The participants were randomly assigned to one of three mentioned between-participant conditions in a way that has a balance across the conditions. After giving their consent, participants needed to take part in three phases of the experiment:

Phase 1 - Pre-experiment Questionnaire: (i) Asking demographic information (age, gender, education, and field of study) (ii) Explaining details of the experiment (iii) Showing a video of collaboration (iv) Asking participants' prior experience with robots (v) Asking participants' prior trust in the robot (Muir's questionnaire [17] with a seven-point Likert scale - Q1-4 in Table 1)

Phase 2 - Main Experiment: (i) Connecting and giving remote control via Zoom (ii) Explaining how to work with the simulation environment (iii) Starting the main experiment (iv) Asking verbally to explain their strategy

Phase 3 - Post-experiment: (i) Asking Q1-4 to measure participants trust in robot after the experiment (ii) Asking Q5-8 to assess participants perception of the robot

4 RESULTS AND DISCUSSION

To show how the participants' trust changed after the experiment, we took the average of the question 1-4, both in the pre- and post-study questionnaires, and calculated their differences (see Figure 2). The results show that for the first and the third conditions participants' trust in the robot increased post-experiment. Figure 3 shows the average of participants answers to Q5-8. As expected, participants gave a high score to question 5 of Table 1, because the robot

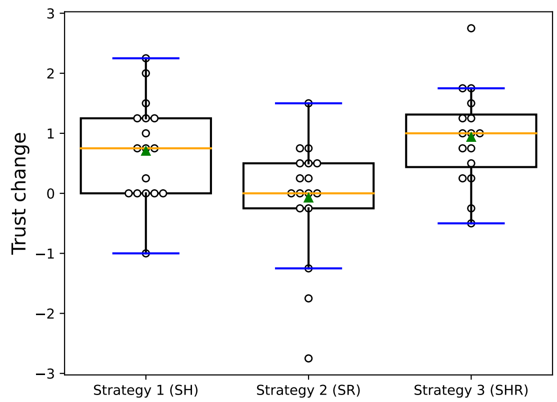

Figure 2: Change trust before and after the experiment

,

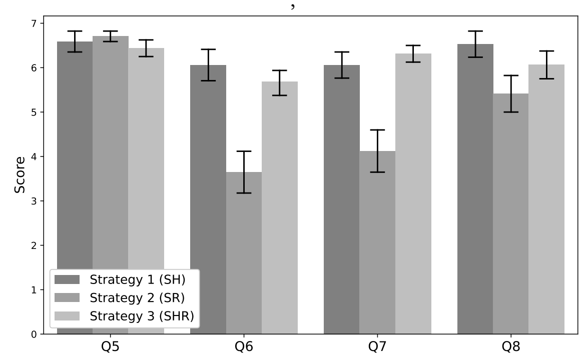

Figure 3: Average answers to the questions in Table 1, Q5-8

can always successfully manipulate objects. Since the questionnaire data does not have a normal distribution, we used a Kruskal Wallis test to evaluate the statistical difference between the three different strategies. Next, Dunn's test, a non-parametric pairwise multiple comparisons procedure, was used as a Post-hoc test. Table 2 shows the results of these tests indicating a significance between strategies 1 (SH) and 2 (SR) and strategies 2 (SH) and 3 (SHR). However, no significant difference was found between strategies 1 and 3.

Participants' strategy: After the experiment, participants were asked to explain the strategy they adopted to accomplish their objective (i.e., finishing the board) and the task (i.e., removing all objects). We also recorded participants' and the robot's actions and analyzed their collaboration with the robot.

Robot strategy 1 (SH): In this case, based on the robot's recorded actions, the robot provides participants with green blocks regardless of their needs. Assuming a self-serving human teammate, we expect that they would give the orange objects to the robot after finishing their own objective. The thick line in Figure 4a shows the number of orange objects that a self-serving human will give the robot. The thin line shows the average number of orange objects that participants gave the robot. Six out of 17 participants followed the self-serving plan, and the others gave at least one orange object before finishing their board. The explanation that participants gave about their strategy revealed two main reasons for doing so: collaborating with the robot to finish the task and preventing the robot from being idle in cases where the robot had no object to move.

Questions	
Q1 - To what extent can the robot's behavior be predicted from moment to moment? Q2 - To what extent can you count on the robot to do its job? Q3 - Overall how much do you trust the robot? Q4 - What degree of faith do you have that the robot will be able to cope with similar situations in the future?	Q5 - The robot is able to manipulate objects successfully and without any failure. Q6 - The robot works efficiently for minimizing collaboration time. Q7 - The robot tends to collaborate. Q8 - I will be happy to collaborate again with this robot in similar future tasks.

Table 1: Questions for humans' trust in the robot (Q1-4) and their perception of the robot (Q5-8)

	SH-SR-SHR		SH-SR	SHR-SR	SHR-SH
	χ^2	p	p (post-hoc test)		
Hypothesis 1a (Q1-4)	12.56	0.0019 (S)	0.01 (S)	0.003 (S)	0.63 (NS)
Hypothesis 1b (Q6)	15.561	0.0004 (S)	0.0004 (S)	0.0122 (S)	0.308 (NS)
Hypothesis 1c (Q7)	12.317	0.002 (S)	0.0088 (S)	0.0043 (S)	0.7008 (NS)
Hypothesis 1d (Q8)	8.168	0.009 (S)	0.013 (S)	0.291 (NS)	0.291 (NS)

Table 2: Statistical difference for hypothesis1 a-d (S: Significant difference, NS: No significant difference).

Robot strategy 2 (SR): The recorded robot's actions show that the robot optimizes its own objective and does not give green objects at first and continues with sorting its own objects. The thick line in Figure 4b shows the number of orange objects that an assistive human teammate (i.e., ones that do not press the skip button when they can provide the robot with orange objects) give to the robot. The thin line shows the average number of orange objects that participants gave to the robot. Six out of 17 participants followed the purely assistive behavior, and the others pressed the skip button at least once despite having the option to give the robot a block. According to participants' answers to our question about their strategy, two of them were annoyed by the robot's behavior, and others wanted to check whether the robot gave them a green object when they pressed the skip button, withholding orange objects.

Robot strategy 3 (SHR): In strategy 3, the robot starts with its own blocks and tries to give green blocks whenever the human needs them. Similar to strategy 1, the thick line in Figure 4c shows the expected number of orange objects that a self-serving human will give to the robot. Three out of 16 participants followed the self-serving behavior, and others gave at least one orange block to the robot although they had the option to place their own blocks on the board. Their explanations for doing so were: (1) Collaborating with the robot to finish the task, (2) Preventing the robot from being idle, (3) Testing if the robot gave them a green object in return.

Prior experience with robot: We also examined how participants' prior experience with robots affected their initial trust in the robot, but the result showed no significant effects.

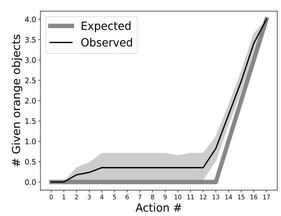

(a) Strategy 1 (SH)

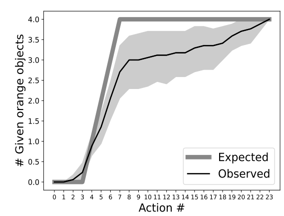

(b) Strategy 2 (SR)

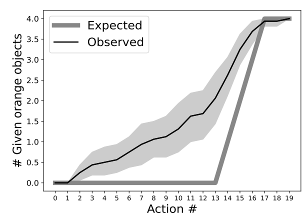

(c) Strategy 3 (SHR)

Figure 4: Expected (thick lines) and observed (thin lines) average number of orange blocks passed to the robot along with their confidence intervals.

5 CONCLUSION

In this paper we considered a single-human single-robot scenario that required the two parties to collaborate to accomplish a given task. The team's task was to remove all blocks from the tables and shared area and put them in the designated bin or assembly board. However, aligned with this task, we assigned the human an objective: to complete their board in the shortest time. Then, we developed three different robot strategies: 1- prioritizing the human (Strategy 1-SH), 2- prioritizing the robot (Strategy 2-SR), and 3- balancing between the human and the robot (Strategy 3-SHR). Through this experiment we sought to find out the extent to which the third strategy can maintain human's trust and satisfaction while optimizing the robot's (or team's) objectives. Based on the results, human perception of the robot was improved by strategies 1 and 3 compared to strategy 2. However, there was no significant difference between strategies 1 and 3. Considering all the above, our findings suggest that, in a collaborating scenario, developing a plan and strategy for the robot that can balance between the human teammate's objective (or policy) and the robot's (or team's) objective can benefit the team without negatively affecting humans' perception of the robot. These results will inform our future work towards improving human-cobot collaboration.

ACKNOWLEDGMENTS

This research was undertaken, in part, thanks to funding from the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Canada 150 Research Chairs Program.

REFERENCES

[1] Arash Ajoudani, Andrea Maria Zanchettin, Serena Ivaldi, Alin Albu-Schäffer, Kazuhiro Kosuge, and Oussama Khatib. 2018. Progress and prospects of the human–robot collaboration. *Autonomous Robots* 42, 5 (2018), 957–975.

[2] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh Rao, and Minoru Asada. 2016. Initiative in robot assistance during collaborative task execution. In *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*. IEEE, 67–74.

[3] Mriganka Biswas and John Murray. 2016. Robots that Refuse to Admit Losing – A Case Study in Game Playing Using Self-Serving Bias in the Humanoid Robot MARC. In *Intelligent Robotics and Applications*. Springer International Publishing, Cham, 538–548.

[4] Karin Bogner, Ulrich Pferschy, Roland Unterberger, and Herwig Zeiner. 2018. Optimised scheduling in human–robot collaboration–a use case in the assembly of printed circuit boards. *International Journal of Production Research* 56, 16 (2018), 5522–5540.

[5] Baptiste Busch, Guilherme Maeda, Yoan Mollard, Marie Demangeat, and Manuel Lopes. 2017. Postural optimization for an ergonomic human-robot interaction. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2778–2785.

[6] Min Chen, Stefanos Nikolaidis, Harold Soh, David Hsu, and Siddhartha Srinivasa. 2018. Planning with trust for human-robot collaboration. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, New York, NY, USA, 307–315.

[7] Filipa Correia, Samuel F Mascarenhas, Samuel Gomes, Patrícia Arriaga, Iolanda Leite, Rui Prada, Francisco S Melo, and Ana Paiva. 2019. Exploring prosociality in human-robot teams. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 143–151.

[8] Sandra Devin and Rachid Alami. 2016. An implemented theory of mind to improve human-robot shared plans execution. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 319–326.

[9] Shirine El Zaatari, Mohamed Marei, Weidong Li, and Zahid Usman. 2019. Cobot programming for collaborative industrial tasks: An overview. *Robotics and Autonomous Systems* 116 (2019), 162–180.

[10] Volker Gabler, Tim Stahl, Gerold Huber, Ozgur Oguz, and Dirk Wollherr. 2017. A game-theoretic approach for adaptive action selection in close proximity human-robot-collaboration. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2897–2903.

[11] Kelsey P Hawkins, Nam Vo, Shray Bansal, and Aaron F Bobick. 2013. Probabilistic human action prediction and wait-sensitive planning for responsive human-robot collaboration. In *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 499–506.

[12] Lars Johannsmeier and Sami Haddadin. 2016. A hierarchical human-robot interaction-planning framework for task allocation in collaborative industrial assembly processes. *IEEE Robotics and Automation Letters* 2, 1 (2016), 41–48.

[13] Poornima Kaniarasu and Aaron M Steinfeld. 2014. Effects of blame on trust in human robot interaction. In *The 23rd IEEE international symposium on robot and human interactive communication*. IEEE, 850–855.

[14] Mahdi Khoramshahi and Aude Billard. 2019. A dynamical system approach to task-adaptation in physical human–robot interaction. *Autonomous Robots* 43, 4 (2019), 927–946.

[15] Theresa Law and Matthias Scheutz. 2021. Trust: Recent concepts and evaluations in human-robot interaction. In *Trust in Human-Robot Interaction*. Elsevier, 27–57.

[16] Yanan Li, Keng Peng Tee, Wei Liang Chan, Rui Yan, Yuanwei Chua, and Dilip Kumar Limbu. 2015. Role adaptation of human and robot in collaborative tasks. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 5602–5607.

[17] Bonnie M Muir and Neville Moray. 1996. Trust in automation. Part II. Experimental studies of trust and human intervention in a process control simulation. *Ergonomics* 39, 3 (1996), 429–460.

[18] Stefanos Nikolaidis, David Hsu, and Siddhartha Srinivasa. 2017. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research* 36, 5-7 (2017), 618–634.

[19] Stefanos Nikolaidis, Ramya Ramakrishnan, Keren Gu, and Julie Shah. 2015. Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 189–196.

[20] Luka Peternel, Nikos Tsagarakis, Darwin Caldwell, and Arash Ajoudani. 2016. Adaptation of robot physical behaviour to human fatigue in human-robot co-manipulation. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, 489–494.

[21] Diede PM Van der Hoorn, Anouk Neerincx, and Maartje MA de Graaf. 2021. " I think you are doing a bad job!" The Effect of Blame Attribution by a Robot in Human-Robot Collaboration. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*. Association for Computing Machinery, New York, NY, USA, 140–148.

[22] Zhikun Wang, Katharina Mülling, Marc Peter Deisenroth, Heni Ben Amor, David Vogt, Bernhard Schölkopf, and Jan Peters. 2013. Probabilistic movement modeling for intention inference in human–robot interaction. *The International Journal of Robotics Research* 32, 7 (2013), 841–858.