



It's a Match: Task Assignment in Human–Robot Collaboration Depends on Mind Perception

Eva Wiese¹ · Patrick P. Weis^{1,2} · Yochanan Bigman³ · Kyra Kapsaskis³ · Kurt Gray³

Accepted: 27 February 2021 / Published online: 26 March 2021
© The Author(s) 2021

Abstract

Robots are becoming more available for workplace collaboration, but many questions remain. Are people actually willing to assign collaborative tasks to robots? And if so, exactly which tasks will they assign to what kinds of robots? Here we leverage psychological theories on person-job fit and mind perception to investigate task assignment in human–robot collaborative work. We propose that people will assign robots to jobs based on their “perceived mind,” and also that people will show predictable social biases in their collaboration decisions. In this study, participants performed an arithmetic (i.e., calculating differences) and a social (i.e., judging emotional states) task, either alone or by collaborating with one of two robots: an emotionally capable robot or an emotionally incapable robot. Decisions to collaborate (i.e., to assign the robots to generate the answer) rates were high across all trials, especially for tasks that participants found challenging (i.e., the arithmetic task). Collaboration was predicted by perceived robot-task fit, such that the emotional robot was assigned the social task. Interestingly, the arithmetic task was assigned more to the emotionally *incapable* robot, despite the emotionally capable robot being equally capable of computation. This is consistent with social biases (e.g., gender bias) in mind perception and person-job fit. The theoretical and practical implications of this work for HRI are being discussed.

Keywords Anthropomorphic robots · Human-likeness · Social robots · Trust · Human–robot interaction

1 Introduction

Collaboration (i.e., offloading parts of one’s own task to others) is a driving force of success and innovation [16]. When a team succeeds, it is because they are able to effectively

manage the unique skills of team members [2]. The importance of optimizing collaborations has given rise to an entire literature of psychological research that documents how people make decisions about joint tasks [9]. As robots become more and more common as work partners, we need to reveal how people choose to rely upon mechanical agents in shared activities [12]. Here, we examine how decisions to outsource tasks to robots are shaped by the perceived fit between the mind of robots and their tasks. This work is crucial as robots and the tasks they are designed for continue to diversify.

Industrial-organizational psychology has long revealed the importance of person-job fit [18]. Selecting the right person for a job means selecting someone whose qualifications correctly match the job demands. These demands can be obvious and physical; for instance, being a firefighter requires physical stamina and the ability to endure extreme situations. However, these demands can also be more social in nature; for instance, being a successful salesperson requires having a likable personality and high emotional intelligence. Researchers have examined how people make decisions of person-job fit [4] and detailed the biases that distort these decisions. For example, one large body

Eva Wiese and Patrick P. Weis authors share first authorship.

✉ Patrick P. Weis
patrick.weis@uni-ulm.de; pweis@gmu.edu

Eva Wiese
ewiese@gmu.edu

Yochanan Bigman
ybigman@unc.email.com

Kyra Kapsaskis
kyrakaps@live.unc.edu

Kurt Gray
kgray@unc.email.com

- ¹ George Mason University, Fairfax, VA, USA
- ² Institute of Psychology and Education, Ulm University, Albert-Einstein-Allee 47, 89069 Ulm, Germany
- ³ University of North Carolina, Chapel Hill, NC, USA

of research finds that people are biased by the gender of applicants, believing that men are better suited for analytic positions and women are better suited for socio-emotional positions [14]. Other work finds a similar bias with respect to combat veterans, seeing military service as preparing veterans for analytic but not socio-emotional jobs [17].

Here we examine people's decisions about robot-job fit, investigating how people delegate tasks to different kinds of artificial agents. It is clear that people implicitly recognize the importance of robot-job fit when it concerns physical characteristics. For example, to construct a car on an assembly line, you need a strong and dexterous robot capable of moving large sheets of metal with precision. However, this kind of industrial robot would not be optimal for interacting with children—and its strength and speed may even prove dangerous.

Just as people use the mind of people to evaluate *person-job* fit, we suggest that they also use the robot's perceived mental capabilities when determining *robot-job* fit. While robots lack a humanlike mind, they are capable of mimicking many mental abilities, including not only advanced computation and memory, but also some rudimentary socio-emotional abilities. Prior research shows that people do perceive some kind of mind in robots, often along the two-dimensional framework, where perceptions of agency (i.e., capacity to think and plan) are separable from perceptions of experience (i.e., capacity to feel and respond emotionally; [8]). Past work reveals that these dimensions of mind perception predict judgments of person-job fit [17]. Accordingly, we suspect that perceptions of agency and experience towards robots will predict what tasks people are willing to assign them. We examine how people coordinate collaboration with two robots, both of which are described as capable of agency but differing in their capacity for experience, with one robot described as being capable of feeling emotions, and the other robot as lacking this ability. We are interested in testing people's willingness to assign two different kinds of tasks to these robots: an arithmetic task requiring computation, and a social task requiring emotional intelligence.

Three important theoretical questions are addressed with the study. The *first question* is whether people will trust the socio-emotional robot to complete the emotional task. It seems obvious that people should select the high-experience robot for the socio-emotional task. This prediction would demonstrate a basic understanding of robot-job fit in the assignment of robot tasks; it would also demonstrate that people legitimately believe that robots possess the ability for some socio-emotional processing. The *second question* is the extent to which biases in human person-job fit judgments are revealed in robot-job fit judgments. This question can be answered by examining which robot people select for the arithmetic task. Given that the arithmetic task involves agency one might expect people to be indifferent to the

robots, selecting them equally frequently for this task. However, classic research finds that people often implicitly treat artificial agents similar to how they treat humans [15] which suggests that perceived experience (i.e., capacity to feel and respond emotionally; [8]) factors in as well. Accordingly, we predict that people will demonstrate an intuitive sense that the ability to effectively perform social tasks is inversely related to the ability to effectively perform arithmetic tasks (cf. [11], which is also in line with the importance of perceived experience of an agent for interaction behavior suggested in earlier research [19]). If this assumption about human person-job fit carries forward to robot-job fit, people should select the low-experience robot for the arithmetic task despite the fact that the high-experience robot is equally capable. The *third question* is how rational people are in their decisions for when and how to delegate tasks. People vary in their abilities to solve social and arithmetic tasks. Accordingly, we predict our participants to outsource cognitive processing to a robot when their actual and/or perceived performance is lower than the robot's ability [6].

In this study, we developed a novel collaborative paradigm to test how mind perception influences robot-job fit. This paradigm varied robot characteristics (high agency, low experience vs. high agency, high experience) and job requirements (socio-emotional vs. computational), and we measured preference decisions. Participants were told that both robots were equally good at both tasks, and we predict that the binary nature of this task—choose robot A or robot B—will reveal any subtle robot-job fit biases.

2 Methods and Materials

Participants engaged in two different tasks: arithmetic and social. In the *arithmetic task*, participants were presented with random arrangements of black and gray dots and were asked to count the dots and report the numerical difference between black and gray dots. In the *social task*, participants were presented with images of the eye region of different human faces and were asked to judge which emotion or experience is depicted. Participants were given a choice between two general response options: (1) answer the question *without help* by selecting one of four answer options depicted on the screen, or (2) *offload* the cognitive processing to one of two robot agents that were represented via images and names on the screen and let them choose the response. Participants were told that selecting one of the robots as response option would lead to the question being answered by that robot's algorithm. They were also told that both robot agents had experience with both tasks (i.e., have solved the tasks before and were equally successful in doing so). In addition, one robot was described as having high emotional capacities (i.e., high in experience) and the

other robot as having low emotional capacities (i.e., low in experience).

2.1 Participants

A total of 157 participants were recruited via Amazon Mechanical Turk (www.mturk.com), of which 143 completed the study. One participant was excluded because of taking more than 45 min for a study designed to take 20 min, resulting in a total sample size of 142 participants (69 female, mean age: 35.3, age range: 18–68; median duration: 18 min, duration range: 13–21 min). All participants reported to be fluent speakers of English and all except for five participants reported English to be their native language. All participants gave informed consent prior to participating and received \$2 for their participation. Participants completed the study online on their own devices. The experiment was presented using the psychological testing software *Inquisit* (www.millisecond.com). Stimulus presentation scaled with the size of the participant's individual screen.

2.2 Stimuli and Task

2.2.1 Robot Agents

Two different robot images were used throughout the experiment: one from the robot EDDIE (e.g., [13]) and one from the robot KISMET [3]. Both robots were of mechanistic appearance and were obtained based on a search for “mechanistic robot” using Google. Mechanistic instead of humanoid robots were chosen to avoid attributions of human-likeness. The robots were named TRF/20 and R-Tec 1, with the name-robot assignment being counterbalanced across participants. One robot was described as capable of experiencing complex emotional states (i.e., low experience: “emotional”), the other described as incapable of experiencing complex emotional states (i.e., low experience: “emotionless”). Robot-description assignment was counterbalanced, and the robots were described as equally capable of solving both tasks.

2.2.2 Social and Arithmetic Task

In total, 72 stimuli were created: 36 for the arithmetic and 36 for the social task. Stimuli for the social task were taken from the Reading the Mind in the Eyes Task [1]; stimuli for the arithmetic task were created using image manipulation software. The dot arrays contained either nineteen or twenty dots, with nine possible numerical differences of black relative to gray dots: -4 , -3 , -2 , -1 , 0 , 1 , 2 , 3 , and 4 .

For both tasks, participants had two different response options: (1) answer the question without any help by selecting one of four response options depicted on the screen, or (2) offload decision making onto one of two robot agents

that were represented via images and names on the screen. Participants were told that selecting (2) would lead to the question being answered by the respective robot's algorithm. The answer options of the social task could arguably include infrequent words (e.g., despondent, incredulous), participants were able to press a small “show dictionary” button at the bottom of the screen that would explain the answer options (this option was used infrequently, on 0.3% of all social task trials). Task trials are illustrated at <https://youtu.be/ROa3BwBxXDA>.

2.3 Procedure

After participants had given informed consent, general instructions concerning the arithmetic and the social task were given. Participants first completed one practice trial for each task type *without help*, were then introduced to *offloading*, and completed one practice trial per task type using offloading. Different robot agents were used for the practice trials than for the experimental trials (i.e., not EDDIE or KISMET).

After participants were familiar with the general mechanics of the task, they were introduced to the robots: TRF/20 and R-Tec 1. To make sure that participants read the instructions carefully, they were asked to answer multiple-choice questions regarding the robots' emotional capacities and autonomy (“What is *robot X* capable of?”).

If participants answered at least one question incorrectly, they were asked to read the instructions again and the procedure was repeated until they were able to answer both questions. Participants were then asked to rate their own, as well as the robots' assumed proficiency in solving both tasks on a visual analogue scale ranging from “very unproficient” to “very proficient” (the left-hand side was coded as 0, the right-hand side as 100). Before the beginning of the first experimental block, participants were reminded of the “offloading” option and that although the robots had prior experience with both tasks they were not necessarily always correct (please note that both robots answered the social and the analytical task with the same accuracy of 67%). They were also told that their task was to score as many correct answers as possible—whether this was accomplished *without help* or through *offloading*. The experiment consisted of 36 trials of the social and 36 trials of the arithmetic task. To incentivize participants to compare the task fit of both robots, they had to rely on the robots (*offloading*) in six arithmetic and six social trials. Specifically, participants were free to choose between both robots but were not allowed to choose an answer option on their own. These twelve trials were excluded from analysis. After successful completion of the 72 trials, participants completed a short demographic survey and performed a second round of proficiency ratings.

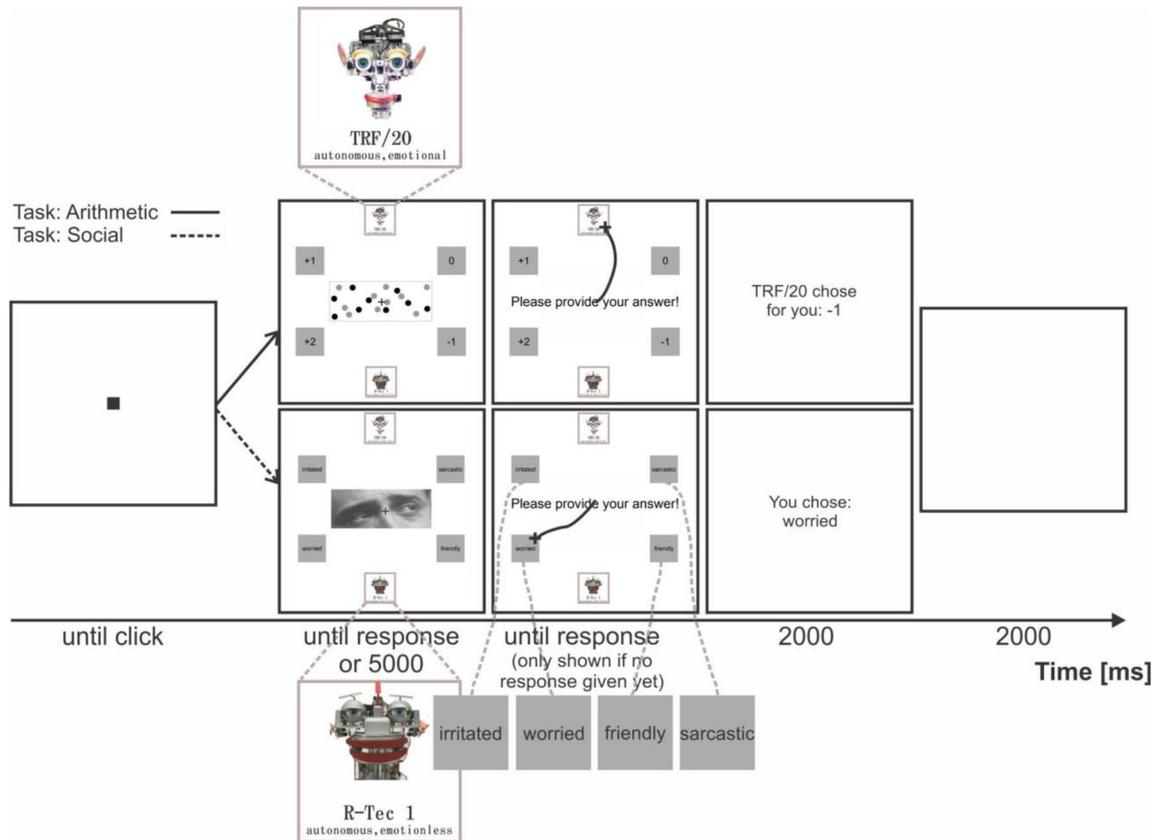


Fig. 1 Trial Sequence. At the beginning of a trial, participants had to click a square to center the mouse cursor. Then, the task-related stimulus (dots for the arithmetic task or eyes for the social task) and the answer options (six squares: four responses and two robots) were shown. If the participant took longer than five seconds to respond, the task-related stimulus disappeared. The five-second window was chosen to allow for informed guesses above chance level for the arithme-

tic task but to simultaneously make highly certain answers unlikely, a pattern that was supposed to match the decision process for the social task. After choosing a response, feedback was provided, followed by a blank screen. Solid black lines with cross at the end illustrate mouse cursor movements. Task trials are illustrated at <https://youtu.be/ROa3BwBxXDA>

Finally, participants were debriefed and thanked for their participation.

The sequence of events throughout the experiment is shown in Fig. 1. On a given trial, participants first clicked on a square in the center of their screen, which was followed by the presentation of the task stimulus, as well as the six response boxes: four for the *without help* condition and two for the *offloading* condition. Response options for the *without help* condition were placed left and right of the task stimulus; response options for the *offloading* condition were presented above and below the task stimulus. Participants then gave their response by selecting one out of the six response options via mouse click.

2.4 Analysis

Three main analyses were conducted: First, we assessed whether task type (arithmetic vs. social) had an impact on offloading behaviors (i.e., number of trials on which

participants let one of the robots respond) and/or accuracy during “without help” trials (i.e., percent correct responses) using a series of paired t-tests.

Second, we examined whether task type (arithmetic vs. social) had an impact on agent preferences (i.e., difference in frequency with which the emotional robot was chosen over the emotionless robot) on “offloading trials” using paired t-tests. Since there were 30 trials per task type (the incentivizing trials were not included in this analysis), a score of **+30** for a given task type would indicate that the emotional robot was preferred on all trials over the emotionless robot; a score of **-30** for a given task type would indicate that the emotionless robot was preferred on all trials over the emotional robot. Thus, when interpreting the preference score, the more positive the value, the more strongly participants preferred the emotional robot, the more negative the value, the more strongly participants preferred the emotionless robot.

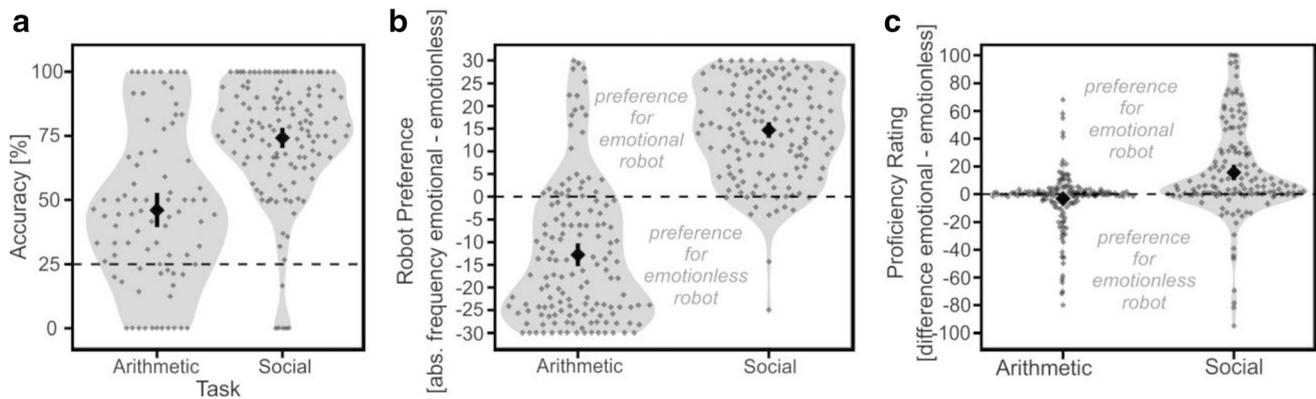


Fig. 2 Results. Accuracy was significantly higher for the social than the arithmetic task (a; dotted line represents chance level). The emotionless robot was preferred for the arithmetic and the emotional robot was preferred for the social task (b). Participants judged the

emotional robot to be more proficient in the social task (c: Social). However, the reverse was not true for the arithmetic task (c: Arithmetic). Error bars represent 95% CIs

Third, we examined whether task type (social vs. arithmetic) had an impact on self- and robot-proficiency ratings prior to engaging in the task (i.e. pre-proficiency ratings) using a series of paired t-tests. Proficiency scores for self- and robot-ratings ranged between +100 and –100. For the *self assessment*, +100 would indicate that the participant considered themselves as maximally proficient on the social task and minimally proficient on the arithmetic task; –100 would indicate that the participant considered themselves as maximally proficient on the arithmetic task and minimally proficient on the social task. That means, when interpreting the self-proficiency ratings, positive scores represent participants who considered themselves as more proficient in performing the social than the arithmetic task, and negative scores represent participants who considered themselves as more proficient in performing the arithmetic than the social task; a score of 0 represents participants who considered themselves as equally proficient for the social and the arithmetic task. Analogously, for the *robot assessment*, +100 means that the participant perceived the emotional robot as maximally proficient and the emotionless as minimally proficient for a given task; –100 means that the participant perceived the emotionless robot as maximally proficient and the emotional robot as minimally proficient for a given task. Thus, when interpreting the robot-proficiency ratings, positive scores indicate a higher perceived proficiency of the emotional than the emotionless robot, and negative scores indicate a higher perceived proficiency of the emotionless than the emotional robot; a score of 0 represents equal perceived proficiency of the emotional and the emotionless robot for a given task.

For all figures, black diamonds represent the mean, error bars 95% confidence intervals (CIs), gray diamonds raw data points, and gray shapes the distribution of the raw data.

3 Results¹

We conducted a series of analyses to answer the three questions of interest. The first analysis (*offloading behaviors*) provides an oversight over our participants' collaboration behavior and showed that participants offloaded the task to one of the two robot agents in 67% of all trials, with a significantly higher offloading rate for arithmetic task trials ($M = 77\%$, $SD = 30\%$) than social task trials ($M = 57\%$, $SD = 32\%$; $t(141) = 8.2$, $p < 0.0001$). Performance on *without help* social task trials ($M = 74\%$, $SD = 23\%$) was higher than on arithmetic task trials ($M = 46\%$, $SD = 30\%$; $t(80) = 8.3$, $p < 0.0001$), which indicates different difficulty levels for the two task types; see Fig. 2a.

The second analysis (*agent preference*) was conducted to answer the first—do people prefer the high-experience robot for the emotional task?—as well as the second—do people prefer the low-experience robot for the arithmetic task despite the fact that the high-experience robot is equally capable?—question. It showed that task type affected which robot was chosen on *offloading* trials: the “emotional” high-experience robot was chosen more often than the “emotionless” low-experience robot for social task trials ($M = 14.7$, $SD = 10.7$); the “emotionless” robot was chosen significantly more often than the “emotional” robot for arithmetic task trials ($M = -12.8$, $SD = 15.5$; $t(141) = 15.2$, $p < 0.0001$); see Fig. 2b. This pattern was mirrored on incentivizing trials, such that participants relied more on the “emotional” versus “emotionless” robot for social task trials (5.3 vs. 0.7 out of 6 trials), and less often on the “emotional” versus

¹ The associated R analysis script and data files can be freely accessed online through the Open Science Framework at <https://osf.io/9y47n/>

“emotionless” robot for arithmetic task trials (1.4 vs. 4.6 out of 6 trials); $t(141) = 18.3, p < 0.0001$.

The third analysis (*proficiency assessments*) was conducted to answer the third question: how rational are people in their decisions regarding when and how to delegate tasks? It showed that participants rated their own proficiency for the social task ($M = 68.8, SD = 22.1$) as being higher than for the arithmetic task ($M = 39.0, SD = 24.4; t(141) = 12.3, p < 0.0001$). The robot-proficiency ratings were modulated by task type ($t(141) = 5.2, p < 0.0001$), which indicates that participants considered the robots as not equally proficient in both tasks. However, unlike the agent preference results, differences in robot-proficiency ratings were driven by significant differences in social ($M_{\text{emotional}} = 75.2, SD_{\text{emotional}} = 23.2; M_{\text{emotionless}} = 59.6, SD_{\text{emotionless}} = 29.7$) rather than arithmetic ($M_{\text{emotional}} = 76.3, SD_{\text{emotional}} = 21.5; M_{\text{emotionless}} = 79.6, SD_{\text{emotionless}} = 16.6$) proficiency: the “emotional” robot was considered as significantly more proficient than the “emotionless” robot for the social task ($M(\Delta) = 15.6$; significantly different from 0: $t(141) = 5.4, p < 0.0001$); both robots were considered as equally proficient in the arithmetic task ($M(\Delta) = 3.2$; not significantly different from 0: $t(141) = 1.8, p = 0.068$); see Fig. 2c.

4 Discussion

The aim of the current study was to examine whether humans prefer robots with high emotional capacities to robots with low emotional capacities, and whether robot preferences were modulated by the to-be-performed task and one’s own as well as the robots’ perceived task proficiency. First, we examined whether participants would trust a robot that was described as emotional to complete a social task although robots are not generally perceived as agents with emotional capacities. Second, we examined whether similar biases as with human-task fit, would also exist when assessing robot-task fit. Third, we investigated whether humans would be rational when offloading task responsibilities to robots such that offloading preferences are informed by whether the robot is considered as particularly proficient in performing the task at hand.

The results show that participants offloaded the task to the robots in the vast majority of trials rather than responding themselves—a pattern that was more pronounced for the arithmetic than the social task. The results also show that agent preferences were driven by robot-task fit, such that the emotional robot was chosen more often to respond to the social task and the emotionless robot was chosen more often to respond to the arithmetic task. Interestingly, while the preference for the emotional over the emotionless robot for the social task was also reflected in higher proficiency ratings for the emotional robot for the social

task, the preference for the emotionless over the emotional robot for the arithmetic task was not reflected in higher proficiency ratings for the emotionless robot for the arithmetic task (i.e., both robots are seen as equally competent to perform the task). The findings indicate that participants are generally willing to offload tasks to robots, and that they calibrate their offloading behaviors based on robot-task fit. Specifically, participants show no reluctance in offloading responses during social task trials to the emotional robot, despite the fact that robots are traditionally not considered as being emotionally capable. The results also suggest that whether a robot agent is considered to be a good fit for a given task is subject to biases, similar to those observed when assessing a human’s task fit: specifically, the results show that although the emotional robot is seen as equally proficient for the arithmetic task as the emotionless robot, it is chosen significantly less often for this type of task than the emotionless robot, indicating that high emotional capacities seem to implicitly (i.e., reflected in offloading behaviors) but not explicitly (i.e., reflected in subjective assessment of proficiency) lower expectations regarding the same agent’s capacities on a task that requires supposedly complimentary skills. A similar bias is observed when it comes to the rationality of participants’ offloading behaviors: specifically, offloading is not only observed for task types that are difficult to perform for the participant (i.e., arithmetic task), it is also observed in more than 50% of all social task trials, which participants had an easy time performing, and considered themselves as being proficient in.

The results highlight that participants have intuitive understanding that robot agents with different capabilities are differentially fit to perform different tasks. Even more remarkably, this preference pattern based on robot-task fit was apparent although the robots were introduced as being equally capable to perform the social and arithmetic task. It also suggests participants legitimately believed that robots possess socio-emotional capacities in the context of a task that requires collaboration. This is a somewhat surprising finding, given that previous research has shown that participants are reluctant to ascribe socio-emotional capabilities to robots when being explicitly asked using subjective ratings [7]. These findings add to the literature on mind perception and human–robot interaction by showing that the context in which a construct is examined strongly determines the outcome: once an agent’s abilities become relevant because collaboration is required, task responsibilities are assigned based on the agent’s believed capabilities and the agent is trusted to make a correct decision.

The current findings also indicate that participants are subject to biases when selecting robot agents for a collaborative task: it is remarkable that although both robot agents are instructed as being agentic, which suggests that they are both equally capable of performing the arithmetic

task, participants preferred the emotionless to the emotional robot for the arithmetic task. This is even more surprising given that when being explicitly asked to rate the proficiency of the robots for the arithmetic task, participants state that they consider them both as equally capable of performing the task. This indicates that implicit biases, similar to gender biases when selecting human employees, seems to be inversely related to the capacity for experience and reduce the likelihood that a perfectly capable agent is selected for a task..

Participants show no reluctance to offload task responsibilities to robot agents that were believed to have prior experience with a set of different tasks. The general offloading tendency was even more pronounced for task types that participants perceived themselves as not being particularly proficient in and that they had difficulties to perform. Specifically, participants offloaded task performance significantly more often to the robot agent during arithmetic than social task trials—a pattern that was matched by lower self-proficiency ratings and accuracy on “without help” trials for the arithmetic versus social task. In addition to participants’ perceived and actual task proficiency, a robot’s subjective fit for a task seems to play a role for offloading: the results show that participants chose the robot with high emotional capabilities significantly more often for the social task and the robot with low emotional capabilities significantly more often for the arithmetic task. The observation that nonhuman helpers are picked based on their perceived task expertise is in line with Dzindolet et al. [5] who showed that participants have preconceived notions regarding the utility of automated aids for different task types. The finding that human and machine aids are selected for advice based on their stereotypical features, is in line with [11] who showed that when seeking advice, participants prefer human agents to machine agents for social tasks and machine agents to human agents for arithmetic tasks. Taken together, these results show that offloading preferences in human–robot interaction seem to be influenced by proficiency considerations both on the human and the robot side, and that task performance is outsourced to robots specifically when one’s own capabilities for the task are low and the robot seems to be generally fit to perform the task.

Overall, the results demonstrate that participants have an implicit understanding of the importance of robot-job fit and allocate tasks to robot agents to (1) generally reduce their workload, (2) compensate for their own perceived and/or actual difficulties with a given task, and (3) maximize the chance for a correct response if the robot is seen as being competent. This has important implications for the field of human–robot interaction. First, the findings highlight the importance of objective measures to examine the impact of pre-existing beliefs on social-interactive processes. Specifically, although participants do not seem to *explicitly* ascribe

socio-emotional capacities to robots when being assessed via subjective ratings [7]) participants *implicitly* consider robots as emotionally capable when offloading socio-emotional tasks in a collaborative setting that requires trust. This shows that the assessment of social processes in HRI requires paradigms that plausibly simulate an interactive environment and measure participants’ responses to robots objectively [20]. Second, the present results show that robots can be perceived as entities with emotional capacities that can be trusted to perform socio-emotional tasks. The challenge for HRI is to determine how robots can activate such perceptions in natural settings, where explicit instructions are not feasible. This suggests that robots need to be designed to be perceived as agents with socio-emotional capacities in order to be sufficiently trusted on social tasks [10]. Third, due to the bias that being perceived as being particularly good at performing one task potentially excludes a robot from being perceived as being good at performing a task requiring a different skill set, the question arises whether it is feasible to aim for a uniform robot design suitable to serve social and computational contexts. The current study suggests that although the emotional robot was perceived to be equally capable to perform the arithmetic task as the emotionless robot, it was not trusted to the same extent as the emotionless robot and selected less often for the arithmetic task than the emotionless robot.

5 Conclusion

The rise of autonomous robots leaves no doubt that soon people will be sharing their workplaces with machines, but whether and how humans are willing to collaborate with them by offloading parts of their own task to the robot is not sufficiently understood. This study shows that humans are indeed willing to offload task assignments in human–robot collaboration, and that the degree of collaboration depends on the kind of mental capabilities (here: high vs. low emotional experience) ascribed to a robot. Similar to biases observed in studies on human person-job fit, this attribution process is subject to implicit biases, which can negatively impact the effectiveness of human–robot collaboration. Research in HRI needs to address the issue of implicit biases towards robots in collaborative work environments and identify effective interventions to mitigate these issues.

Funding Open Access funding enabled and organized by Projekt DEAL. YB acknowledges an NSF-SBE Post-doctoral fellowship, and KG acknowledges the Charles Koch Foundation.

Data Availability Data, analysis scripts, and study materials are available to other researchers using the OSF repository at <https://osf.io/9y47n/>.

Declarations

Conflict of interest The authors declare that they have no affiliations with or involvement in any organization or entity with any financial or non-financial interest in the subject matter or materials discussed in this manuscript.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baron-Cohen S, Wheelwright S, Hill J, Raste Y, Plumb I (2001) The "Reading the Mind in the Eyes" Test revised version: A study with normal adults, and adults with Asperger syndrome or high-functioning autism. *J Child Psychol Psychiatry Allied Discipl* 42(2):241–251
- Belbin RM (2012) *Team roles at work*. Routledge, London
- Breazeal, C., & Scassellati, B. (1999). How to build robots that make friends and influence people. *Proceedings of the 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS99)*, 858–863.
- Chatman JA (1989) Matching people and organizations: selection and socialization in public accounting firms. *Acad Manag Proc* 1989(1):199–203. <https://doi.org/10.5465/ambpp.1989.4980837>
- Dzindolet MT, Pierce LG, Beck HP, Dawe LA (2002) The perceived utility of human and automated aids in a visual detection task. *Hum Factors* 44(1):79–94. <https://doi.org/10.1518/0018720024494856>
- Fitzsimons GM, Sackett E, Finkel EJ (2016) Transactive goal dynamics theory: a relational goals perspective on work teams and leadership. *Res Organ Behav* 36:135–155. <https://doi.org/10.1016/j.riob.2016.11.006>
- Gray HM, Gray K, Wegner DM (2007) Dimensions of mind perception. *Science* 315(5812):619–619
- Gray K, Wegner DM (2012) Feeling robots and human zombies: mind perception and the uncanny valley. *Cognition* 125(1):125–130. <https://doi.org/10.1016/j.cognition.2012.06.007>
- Gruenfeld DH, Mannix EA, Williams KY, Neale MA (1996) Group composition and decision making: how member familiarity and information distribution affect process and performance. *Organ Behav Hum Decis Process* 67(1):1–15. <https://doi.org/10.1006/obhd.1996.0061>
- Hertz N, Wiese E (2018) Under pressure: examining social conformity with computer and robot groups. *Human Factors*. <https://doi.org/10.1177/0018720818788473>
- Hertz N, Wiese E (2019) Good advice is beyond all price, but what if it comes from a machine? *J Exp Psychol Appl* 25(3):386
- Hinds PJ, Roberts TL, Jones H (2004) Whose job is it anyway? A study of human-robot interaction in a collaborative task. *Hum-Comput Interact* 19(1):151–181. https://doi.org/10.1207/s15327051hci1901&2_7
- Kühnlentz K, Sosnowski S, Buss M (2010) Impact of animal-like features on emotion expression of robot head eddie. *Adv Robot* 24(8–9):1239–1255
- Moss-Racusin CA, Dovidio JF, Brescoll VL, Graham MJ, Handelsman J (2012) Science faculty's subtle gender biases favor male students. *Proc Natl Acad Sci* 109(41):16474–16479. <https://doi.org/10.1073/pnas.1211286109>
- Nass C, Moon Y (2000) Machines and mindlessness: social responses to computers. *J Soc Issues* 56(1):81–103. <https://doi.org/10.1111/0022-4537.00153>
- Nieto MJ, Santamaría L (2007) The importance of diverse collaborative networks for the novelty of product innovation. *Technovation* 27(6):367–377. <https://doi.org/10.1016/j.technovation.2006.10.001>
- Shepherd S, Kay AC, Gray K (2019) Military veterans are morally typecast as agentic but unfeeling: implications for veteran employment. *Organ Behav Hum Decis Process* 153:75–88. <https://doi.org/10.1016/j.obhdp.2019.06.003>
- Walsh WB, Holland JL (1992) A theory of personality types and work environments. In: Walsh WB, Craik KH & Price RH (eds) *Person-environment psychology: models and perspectives*. Erlbaum, Hillsdale, NJ, pp 35–78
- Waytz A, Gray K, Epley N, Wegner DM (2010) Causes and consequences of mind perception. *Trends Cogn Sci* 14(8):383–388
- Wiese E, Metta G, Wykowska A (2017) Robots as intentional agents: using neuroscientific methods to make robots appear more social. *Front Psychol*. <https://doi.org/10.3389/fpsyg.2017.01663>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.