2020 29th IEEE International Conference on
Robot and Human Interactive Communication (RO-MAN)
August 31 - September 4, 2020 - Virtual Conference

Adaptive Leader-Follower Behavior in Human-Robot Collaboration

Emma M. van Zoelen, Emilia I. Barakova, and Matthias Rauterberg

Abstract— **As developments in artificial intelligence and robotics progress, more tasks arise in which humans and robots need to collaborate. With changing levels of complementarity in their capabilities, leadership roles will constantly shift. The research presented explores how people adapt their behavior to initiate or accommodate continuous leadership shifts in human-robot collaboration and how this influences trust and understanding. We conducted an experiment in which participants were confronted with seemingly conflicting interests between robot and human in a collaborative task. This was embedded in a physical navigation task with a robot on a leash, inspired by the interaction between guide dogs and blind people. Explicit and implicit feedback factors from the task and the robot partner proved to trigger humans to reconsider when to lead and when to follow, while the outcome of this differed across participants. Overall the participants evaluated the collaboration more positively over time, while participants who took the lead more often valued the collaboration more negatively than other participants.**

I. Introduction

As robots gain more autonomy, they will take over responsibilities of people. As these robots become more ubiquitous, more and more tasks will arise in which humans and robots have to work together to make optimal use of the qualities of both. This has been studied for quite a while, especially in safety-critical contexts such as the military [1], space applications [2], [3] and search and rescue problems [4], [5]. More recently, the question of whether intelligent robots should be considered partners rather than merely tools has been growing as a more general topic in the human-computer interaction, social robotics and design community [6], [7].

When the capabilities of collaborating humans and robots are complementary, situations might arise in which the intentions of the human and the robot are seemingly conflicting, because the partners do not fully understand the mental processes happening in the mind of the other. Such situations may cause shifts in leadership roles, creating a necessity to balance leading and following by adapting to each other in the collaborative process. After collaborating for a while, these shifts might become more natural as both systems develop an understanding of the other's behavior in the context of the task by adapting their mental models. Following this, the main purpose of the work presented in this paper is to explore and observe how humans adapt their behavior to initiate or accommodate shifting leadership roles and deal with conflicting intentions in human-robot collaborative tasks.

Recent research on joint action coordination (e.g. [8]) and human-robot collaboration (e.g. [9]) has shown that balancing leading and following is an important mechanism that enables coordination and helps to increase the human's understanding of the robot. It can be said that such a process helps to establish and maintain common ground, one of the main aspects necessary for enabling collaboration between humans and intelligent robots [10]. This might also be called mutual understanding, meaning that both parties are able to predict and/or explain the other's actions, leading to trust and eventually smooth collaboration [11]. Moreover, it has been suggested that both parties need to adapt their behavior to the other to achieve mutual understanding [12]. For the study we present in this paper, we have developed an experimental paradigm that facilitates the investigation of behavioral adaptation and emergent role switching in leader-follower situations, as well as how it influences the understanding and trust between a human and a robot.

Using this paradigm, we conducted an experiment in which participants performed a task collaboratively with a robot operated via Wizard of Oz (WoZ). During the task, participants were constantly confronted with seemingly conflicting intentions of the robot. At each such occasion, participants had to decide to either lead or follow the robot, allowing us to observe when they switched between leading and following and how their behavior changed over time. A detailed description of the experimental methods is given in Section III. Results show the different ways in which leader-follower behavior develops over time and how this relates to experienced Collaboration Fluency. Triggers for switching between leading and following roles were identified as well. The results are described further in Section IV.

II. Related Work

A. Human Social Coordination

Studies on human-human collaboration provide a basis for understanding human-robot collaboration, in particular studies on how leader/follower roles are coordinated and how they develop over time. Several studies exist that have looked into 'action signaling' as a way of implicit communication about intentions and (leadership) roles in collaborative activity (e.g. [13]). Some papers evaluate the coordination behaviors when there is an asymmetry in the coordinative task (e.g. predefined leaders and followers, as in [8], [14]). It is shown that both leaders and followers signal their role to the other by means of slight adaptations in their action execution. They thereby make clear how the other should coordinate with them. This also happens when roles are not explicitly defined, but, for example, emerge from the specific abilities of the two actors. These studies, however, deal only with situations in which the (leader and follower) roles are

E. M. van Zoelen was with Eindhoven University of Technology, Eindhoven, 5612AZ, The Netherlands. She is now with Delft University of Technology, Delft, 2628CD, and TNO, Soesterberg, 3769DE The Netherlands. (phone: +31 6 17 345 121; e-mail: emma.vanzoelen@tno.nl).

M. Rauterberg and E I. Barakova are with Eindhoven University of Technology, Eindhoven, 5612AZ The Netherlands. (e-mail: g.w.m.rauterberg@tue.nl, e.i.barakova@tue.nl).

978-1-7281-6075-7/20/$31.00 ©2020 IEEE

static and do not shift or change over time. There is some work on behavioral dynamics of coordination [15] which shows that humans engage in a mutually adaptive process when they need to coordinate. However, the task used in [15] is a very static one and by design the capabilities of the coordinating partners are symmetrical. The experiment presented in this paper used a task in which the capabilities of the partners are asymmetrical (with shifting asymmetries during collaboration), and focuses on the behavioral dynamics of ongoing coordination and the signaling of leader/follower intentions.

B. *Shifting Roles in Human-Robot Collaboration*

In most existing work on Human-Robot Collaboration, the human is generally the leader (also called the master-slave principle). However, to truly study collaboration, partners need to be able to switch roles dynamically, and mutual adaptation between the human and the robot must be possible [9]. Some research work has attempted to let a robot partner dynamically change roles [16], and also to more specifically switch between leader and follower roles based on feedback from a human collaborator [17], [18]. These studies, however, approached role switching mostly as a technical challenge for the robot partner, not focusing on the effects on interaction. While they are empirically evaluated in context with a human, this is either a very small scale evaluation, or the focus is mostly on improvements in performance. In [18] the authors mention the risk of co-adaptation, with which is meant that the human partner adapts to the shifting roles of the robot. This was however not observed. In our work, we focus specifically on this co-adaptation, as we attempt to observe how humans adapt their behavior in the perspective of shifting leadership roles.

C. *Mutual Adaptation*

In human-robot, human-computer and human-machine interaction, a lot of work has been done on letting machines (or computers, robots, agents) adapt to humans to enable smooth collaborations (e.g. [19], [20]). However, those studies generally do not take into account that the human is likely to adapt their behavior as well, and what a mutually adaptive relationship emerging from that may mean for the collaboration. Studies that do take human adaptation into consideration have shortcomings as well. Some do not pay attention to the influence of the behavior of the intelligent machine (e.g. [21]). Others provide the intelligent machine with the ability to deal with an adaptive human, but do not analyze the resulting changes in behavior of the human (e.g. [22]). The purpose of our work is to focus on human adaptation while taking the behavior of the robot into account. We analyze what patterns of behavior emerge and investigate whether they help the human to better understand the intelligent robot's behavior. Ultimately, this can help in creating intelligent robots that can adapt to the human, keeping their partner's adaptation into account as well.

D. *Meaningful Human Control*

The work we present in this paper relates to a relatively new research area on 'Meaningful Human Control', which has arisen as a response to the future possibility of autonomous weapon systems [23], but is currently considered relevant in many other human-robot interaction or collaboration contexts. In this research area, the main objective is to find a way to make sure that humans have the right amount of control over the behavior of an autonomous system, to make it behave as we humans desire without compromising the benefits of its autonomy. The question of who leads and who follows at which point of the task can be translated into a question of who is in control as well.

At present, a lot of the work on Meaningful Human Control (MHC) deals with the discussion of what that term actually means, or what the requirements for MHC are. It is, for example, unclear whether a human should constantly be 'in the loop', or whether humans can also be allowed 'off the loop' [24]. Some studies apply the philosophical concepts to specific applications, such as autonomous vehicles [21] or surgical robots [26]. However, an (empirical) analysis of what MHC entails for humans collaborating with intelligent robots, how their behavior is influenced by the level of control and how it might develop over time is not yet present in these works. While the work presented in this paper does not provide answers for achieving Meaningful Human Control, it does provide an analysis of how leading behavior and as a consequence objective and subjective control develops over time.

III. Method

The described experiment and data collection were approved by the Ethical Review Board of the Industrial Design department at Eindhoven University of Technology (Reference number: ERB2019ID7, approved on 21-11-2019).

A. *Robot Interaction Design*

The design of a task for studying leader-follower behavior in human-robot collaboration was inspired by other interactions between humans and non-humans, in this case between blind people and their guide dogs. This context is an example of situations in which a human and a non-human have complementary capabilities and collaborate towards a common goal (reaching a destination safely). A remote-controlled robot with a leash (Figure 3) and a navigation task were designed. The appearance of the robot was designed without anthropomorphic features on purpose, to allow anyone interacting with it to focus on the interaction. The leash was designed to be the only direct communication channel between the robot and a participant, to ensure specific evaluation of the interaction through the leash without noise of other interaction modalities. The leash interaction allows for subtle and implicit interactions as both the participant and the robot can pull the leash to different degrees. The robot was designed to be large and heavy, to enable it to pull the participant in a direction when required. The robot was controlled remotely by a human Wizard. This ensured that the robot could take its 'own' actions, but could also respond adaptively to its human partner (the participant) when they were pulling the leash.

B. *Experimental Setting*

Participants were told that they had to perform a collaborative task together with an intelligent robot while holding the leash of the robot. They were brought to an experimental football field for robot soccer (Figure 2) and introduced to the robot. The participants were told that they had to move from one point on the field to another point on the other half of the field together with the intelligent robot. They had to score as many points as possible. They were given 60 points at the start of a run, but lost a point for every second it took them to reach the final target location. They were told that they could earn extra points by picking up an unknown number of virtual objects that were hidden in the environment, but only the robot knew where those objects were. A sound would indicate that they had 'picked up an object'. Virtual objects could be near the shortest route from start to finish, but they could also be outside the route, which of course would cost time and points. This setup required the participants to make choices between speed and exploration (and thus between leading and following) in their attempt to achieve a high score. The robot would follow a default route to pick up all virtual objects, unless being pulled away from this route by the human participant.

Participants were first familiarized with the experiment, the field and the robot. They walked from one end of the field to the other with the robot, to give them an indication of the speed of the robot. After that, the first run of the task started. Each participant completed four runs; the locations of the virtual objects were different for every run. Four maps with predefined locations of the virtual objects were constructed for the Wizard (Figure 1). The order of these maps was randomized for each participant.

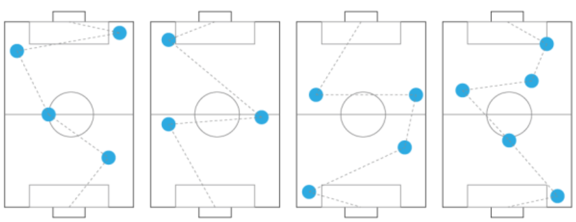

Figure 1. The four predefined maps with the locations of the objects (blue circles), including a line indicating the default route of the robot. The bottom of the field is the starting point.

C. Participants

A total of 18 people participated in the experiment (9 male, 9 female), consisting of students from Eindhoven University of Technology, with an average age of 23 (SD = 3.9). Participants were told that the person with the highest number of points on a single run would receive a gift voucher of €10 to motivate them to perform to the best of their abilities. Before the start of the experiment, the participants gave their consent after reading the consent form that explained all details of the experiment except for the focus of the research and the specific behavior of the robot.

D. Data Collection

Several types of qualitative and quantitative data were collected. While performing the task, a camera placed in a corner of the field recorded the behavior of the participants. After each run, participants were asked to complete a questionnaire on subjective Human-Robot Collaboration Fluency [27], and answered three interview questions:

1. Can you explain the behavior of the robot?
2. What was your strategy for completing the task?
3. How did you experience the collaboration?

Figure 2. The field on which the task was executed. Participants moved from the goal on the left (where the robot is stationed) to the goal on the right.

Figure 3. Two participants interacting with the robot showing a situation with a stretched leash and thus in a leading role (top) and situation with a loose leash and thus a following role (bottom).

E. Data Analysis

The videos were first coded using an open coding process, to get a view on the different behaviors present among participants as well as on events that triggered participants to switch between a more leading and a more following role. From the open codes, a closed coding scheme was developed that contained codes for task events, robot movement, participant movement, leash activity and the participant's location relative to the robot. Codes related to the participant's behavior were marked as leading or following behaviors. All videos were then coded again using a closed coding process with The Observer XT [28], enabling us to visually analyze the different behaviors across runs simultaneously as well as to quantify the amount of leading behaviors present in each run. Intercoder reliability for the duration of sequences with another coder for 5.6% of the data (videos of 4 runs) was found to be 97.55%. Looking at the development of leading behavior across runs for the different categories of participant behavior (leash activity, participant movement and participant location), participants with similar behavior were manually clustered. Similarly, participants were also clustered based on interview answers.

A Repeated Measure ANOVA was conducted on the Collaboration Fluency questionnaire results to see if, within-subject, the scores changed over the different runs. A post-hoc analysis using a Tukey HSD test with Bonferroni correction was done to test between which runs any found differences were significant.

For one of the female participants data was missing in the Collaboration Fluency questionnaire. Therefore, for the analyses relating this measure this person was excluded.

IV. Results

A. Switch Triggers

The open coding process revealed six types of situations which will be referred to as trigger moments in the task. Each such situation typically triggers the participants to reconsider whether they should behave in a more leading or following way. This could, for example, be observed by a clear switch, hesitating behavior or several short switches after one of the described triggers.

The first such moment is at the start of the task, when participants express whether they want to start the task in a more leading or following manner. The other five trigger moments for reconsidering roles are the following:

- Sound indicating a virtual object;
- Leash pull by the robot (when it goes in another direction than the participant);
- The robot deviating from the route that leads to the final goal (without clear leash pull);
- Getting close to the goal;
- The robot standing still.

These five types of trigger moments can be categorized as explicit, implicit or ambiguous feedback, as well as in task or partner feedback as can be seen in Table I. Related to the 'partner feedback' category, it was evident that almost all participants unconsciously followed the robot when it was heading roughly in the direction of the goal. This is clear from the fact that participants' routes often did not lead to the goal in a straight line, but that they followed the robot to the sides of the field (and thus via one of the virtual objects). The robot deviating from the goal route is therefore only really a switch trigger if the deviation from the route is substantial. The implicit task feedback was evident as for most participants, being further along in the task made them change their behavior. While this is, in essence, a gradual change, it was sometimes observable as an immediate change right after people e.g. crossed the middle of the field.

TABLE I. Feedback as Switch Trigger

	Explicit	Implicit	Ambiguous
Task Feedback	Object sound	Getting close to the goal	Robot standing still
Partner Feedback	Leash pull by the robot	Robot deviating from goal direction	Robot standing still

B. Leading Behavior Development

Closed coding made it possible to look at the development of behavior across different runs per participant. We qualitatively looked at the three main categories of behaviors (leash pull, position relative to robot, movement) and created descriptions for the development of the behavior in each category. We then grouped participants with similar developments into a few main groups (Figure 4). Each of these groups was characterized as either a mostly leading, mostly following or balancing behavior (where participants moved towards a more balanced collaboration over time).

For all three different leading behaviors, the following similar developments could be observed:

- Mostly following (following behavior, Figure 4a);
- Start very following, leading in the middle, following at the end (following behavior, Figure 4b);
- Start very following, increase of leading over time (balancing behavior, Figure 4c);
- Start very leading, increase of following over time (balancing behavior, Figure 4d);
- Start very leading, following in the middle, leading at the end (leading behavior, Figure 4e);
- Mostly leading (leading behavior, Figure 4f).

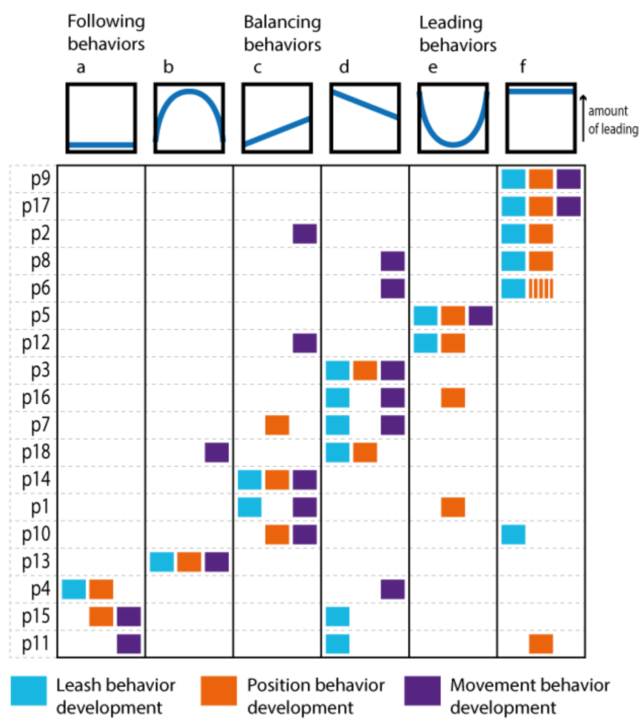

Figure 4. The clustering of the leading behavior development. Categories a to f indicate the different types of development of leading behavior. Blue blocks represent leash behavior development, orange blocks represent relative position behavior development, and purple blocks represent movement behavior development.

There was one participant that did not exactly fit in one of these categories for their relative position development, as they showed a mostly leading behavior that became more fragmented over time (p6). Overall, many participants appear in similar clusters across the different kinds of behavior (e.g. a generally leading behavior for leash activity usually means that there is also a generally leading behavior in the other two categories, or at least a balancing behavior development). In the visualizations of the closed codes this was visible as often (though not always) different leading behaviors appeared simultaneously. An interesting difference between the three kinds of behaviors is that for leash activity, most participants start with a stretched leash. A part of those then continue to balance this out with following behavior, but several others do not show development in this behavior. For relative position, however, most participants remain in a leading position for most of the experiment. Last, the movement of the participants is generally more balancing in the sense that most participants level out their behavior over the runs.

C. Collaboration Fluency

The Repeated Measure ANOVA of the results of the Collaboration Fluency questionnaire for all participants indicated a significant difference between the different runs ($F(3,48) = 6.76$, $p < .001$). The post-hoc analysis showed significant results between runs 1 ($M = 51.11$, $SD = 9.80$) and 3 ($M = 58.86$, $SD = 10.23$), between runs 2 ($M = 54.49$, $SD = 9.19$) and 4 ($M = 62.39$, $SD = 12.97$) and between runs 1 and 4 (see Table V). No differences were found between male and female participants.

It is clear that the subjective Collaboration Fluency grows over time, and that this growth can be seen within three runs. This means that regardless of how people behave, the fact that they interact with the robot by itself makes most people more positive about the collaboration over time.

TABLE II. COLLABORATION FLUENCY RUNS COMPARED

Runs Compared	P-value[a]
1-2	1.000
1-3	0.023*
1-4	<0.001**
2-3	0.623
2-4	0.020*
3-4	1.000

a. Values with * are significant, values with ** are extremely significant.

D. Collaboration Fluency Linked to Behavior

For each leading behavior the total duration relative to the total duration of the task was taken per participant per run. A Pearson correlation test was used to test for correlation between this and the scores for subjective Collaboration Fluency. This yielded the correlation coefficients presented in Table VI.

There is a weak negative correlation between the total duration of a stretched leash and the subjective Collaboration Fluency score. This means that when participants stretch the leash more, they will generally score lower on the questionnaire. A similar but weaker negative correlation can be found between the total duration of the participant moving in the direction of the goal and the subjective Collaboration Fluency score. This suggests that if people portray less (explicit) leading behavior, the Collaboration Fluency is higher. Therefore, when people are less willing to follow the robot, they regard the robot as less cooperative.

TABLE III. CORRELATIONS LEADING BEHAVIOR AND COLLABORATION FLUENCY

	Correlation coefficient	P-value[a]
Stretched leash	-0.344	0.006**
Position in front of robot	-0.113	0.373
Moving in goal direction	-0.290	0.020*

a. Values with * are significant, values with ** are extremely significant.

E. Interview Insights

The interviews were meant to get a view on how participants would subjectively describe the collaboration with the robot in relation to the development of their behavior. For that reason, a description of the development of their answers was made, after which these descriptions were clustered (Table VII). Almost all participants started understanding the behavior of the robot better over the course of the runs, since their explanations of the robot behavior as well as their own strategies became more elaborate and less uncertain. The great majority of the participants felt like the collaboration became better and more balanced over time (which was also clear from the Collaboration Fluency questionnaire). Several participants, however, expressed that the collaboration remained imbalanced until the last run, because one of the partners was clearly more important in achieving the task. While several participants regarded this as a negative aspect, some did not consider this to be a problem or even mentioned that they considered this a positive aspect.

It is interesting to note that the extent to which the collaboration is balanced was expressed so clearly. Overall, participants considered a balanced collaboration in which both partners contributed equally superior to an imbalanced collaboration, even if the latter made it easier to complete the task. Some participants even described how they did their best to find aspects in which they could help the robot with its weaknesses, such as the fact that it is quite slow (participants pull the leash and use verbal encouragement to motivate it to be faster) or the fact that it cannot make sharp turns (one participant picked up the robot and carried it the final distance to the goal because they knew it would take the robot a lot of time to make the turn).

TABLE IV. INTERVIEW CLUSTERS

Subjective description of collaboration	Participants in this cluster
Collaboration becomes better and more balanced	1, 2, 6, 7, 8, 10, 11, 14, 16, 17
Collaboration remains imbalanced, this is regarded negatively	3, 12, 13, 18
Collaboration remains imbalanced, this is regarded positively	4, 9, 15
Robot is more of a tool that needs to be lead	5

F. *Interviews Linked to Behavior*

The interview clusters were visually compared to the behavioral clusters using crosstabulation, to explain why we observed certain behavior. Most participants that report that the collaboration becomes better and more balanced in the interview also have a movement development that becomes more balanced (8/10). For their leash behavior, they either have a balancing development or have a generally leading behavior (5/10, 5/10), which is similar for their position development (3/10, 7/10). For both imbalanced interview categories, the distribution over behavior development styles is quite equal. For the one participant that considered the robot to be more of a tool that needed to be lead, all behavioral measures fall within the leading category.

These numbers suggest that exploring different leading and following behaviors (thus balancing them out) helps participants to empathize with the robot. Also, participants project their own behavior onto the robot; if they themselves behave in a more balancing way, they consider the robot and general task execution to be more balanced and cooperative.

V. DISCUSSION

The research presented in this paper explores how the leading behavior of humans develops over time when they collaborate with a robot partner, focusing on what triggers them to shift leadership from or to the robot. We provide a new perspective on what such collaborative and coordinative behaviors could look like, while explicitly taking into account that human behavior will be continuously influenced by robot behavior in such situations. The insights can be used in designing for different HRC contexts, such as self-driving cars, factory robots, or urban search and rescue. The results contribute to an understanding of:

- Interactions that trigger people to reconsider leadership roles
- How leader/follower behavior changes over time
- The interplay between subjective Collaboration Fluency and shifting leader/follower roles

It is interesting that people's appreciation of the collaboration increases as they do more runs with the robot, no matter what their specific behavioral development is. This result might imply that humans are inclined to trust robots more the longer they collaborate with them, and we might wonder whether that is desirable. In a real-world application, it will be important to ensure that this mechanism does not create overtrust. An interesting direction for future work is to investigate to what extent the behavior of the robot influences this growth of subjective Collaboration Fluency, and what factors are responsible for it.

The finding that people who are inclined to closely follow the robot also have a higher appreciation of their collaboration may have implications for the 'meaningfulness' of Meaningful Human Control. It shows the need to find a way to ensure that people will retake the lead when this is necessary. Again, further research into this effect is necessary as it may inform the development of interaction design that ensures the human remains critical to the behavior of the robot. The various types of explicit and implicit feedback from the partner and the task that were found might be of help in designing a balanced leader-follower relationship between a human and a robot. Using triggers that stimulate humans to reconsider whether they should adopt a leading or following role could prove useful.

People differ substantially from one another in the development of leading and following behavior over time. One explanation for this is that some people take more time to balance out leading and following. Another possibility is that some people just prefer to be more or less in control. The diversity of leadership development will need to be taken into account when designing collaborative interactions between humans and intelligent robots.

The results of the present experiment provide interesting insights, but it is difficult to draw clear conclusions given the diversity of results and the relatively small number of participants. Furthermore, it is important to realize that using a human Wizard who simulates the agency of the robot is an approximation. A robot with AI might behave differently from the Wizard. Finally, the presence of the human Wizard might have had an influence on the behavior of the participants. Further evaluations of the different results and evaluations with intelligent robots are necessary to get a deeper and more reliable insight. The presented study serves as the first step into that understanding by investigating the different interaction mechanisms that build up a shifting leader-follower relationship between humans and intelligent robots in collaborative activity.

VI. CONCLUSION

Understanding how leading and following behavior develops in human-robot collaborations requires experimental work that allows for subtle and implicit interactions. The guide dog-inspired experimental paradigm presented allowed us to study such interactions in detail due to its physical and embodied nature. The results show that people extract implicit and explicit feedback from the behavior of the robot partner as well as from the progression of the task. Throughout the task, participants evaluate their current leading or following role and consider whether or not to hand over or take the lead due to such feedback. Participants appreciated the collaboration with the robot more over time, but overall the appreciation was higher for runs in which participants were inclined to adopt a more following position. It is important that behavioral changes in collaboration, such as the ones found in this study, are taken into account in the design of robots in collaborative tasks. The different types of feedback as switch triggers as well as the leader-follower behavior development clusters can be used for further research into personalization of robot behavior, to build towards future design of positive collaborative experiences between human and machine in which control is mediated in a responsible way.

ACKNOWLEDGMENT

The results presented in this paper are based on work done for the Master Thesis in Industrial Design of Emma M. van Zoelen, where more details about the study and the process leading up to it can be found [29].

REFERENCES

[1] R. Parasuraman, M. Barnes, and K. Cosenzo, "Adaptive Automation for Human-Robot Teaming in Future Command and Control Systems," *The International C2 Journal*, vol. 1, no. Decision Support for Network-Centric Command and Control, pp. 43–68.

[2] J. M. Bradshaw *et al.*, "Adjustable Autonomy and Human-Agent Teamwork in Practice: An Interim Report on Space Applications," in *Agent Autonomy*, vol. 7, H. Hexmoor, C. Castelfranchi, and R. Falcone, Eds. Boston, MA: Kluwer Academic Press, 2003, pp. 243–280.

[3] J. van Diggelen, J. M. Bradshaw, T. Grant, M. Johnson, and M. Neerincx, "Policy-Based Design of Human-Machine Collaboration in Manned Space Missions," in *2009 Third IEEE International Conference on Space Mission Challenges for Information Technology*, Pasadena, CA, USA, Jul. 2009, pp. 376–383, doi: 10.1109/SMC-IT.2009.52.

[4] G. J. M. Kruijff, M. Janıcˇek, S. Keshavdas, B. Larochelle, and H. Zender, "Experience in System Design for Human-Robot Teaming in Urban Search & Rescue," in *8th International Conference on Field and Service Robotics*, Matsushima, Japan, 2012, pp. 1–14.

[5] I. R. Nourbakhsh, K. Sycara, M. Koes, M. Yong, M. Lewis, and S. Burion, "Human-Robot Teaming for Search and Rescue," *IEEE Pervasive Computing*, vol. 4, no. 1, pp. 72–78, Jan. 2005, doi: 10.1109/MPRV.2005.13.

[6] R. K. E. Bellamy, S. Andrist, T. Bickmore, E. F. Churchill, and T. Erickson, "Human-Agent Collaboration: Can an Agent be a Partner?," in *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems - CHI EA '17*, Denver, Colorado, USA, 2017, pp. 1289–1294, doi: 10.1145/3027063.3051138.

[7] A. Ajoudani, A. Zanchettin, S. Ivaldi, A. Albu-Schäffer, K. Kosuge, and O. Khatib, "Progress and Prospects of the Human-Robot Collaboration," *Autonomous Robots*, vol. 42, no. 5, pp. 957–975, 2018, doi: 10.1007/s10514-017-9677-2.

[8] C. Vesper and M. J. Richardson, "Strategic communication and behavioral coupling in asymmetric joint action," *Exp Brain Res*, vol. 232, no. 9, pp. 2945–2956, Sep. 2014, doi: 10.1007/s00221-014-3982-1.

[9] J. Bütepage and D. Kragic, "Human-Robot Collaboration: From Psychology to Social Robotics," *arXiv:1705.10146 [cs]*, May 2017, Accessed: Mar. 03, 2020. [Online]. Available: http://arxiv.org/abs/1705.10146.

[10] G. Klein, D. D. Woods, J. M. Bradshaw, R. R. Hoffman, and P. J. Feltovich, "Ten challenges for making automation a 'team player' in joint human-agent activity," *IEEE Intelligent Systems*, vol. 19, no. 6, pp. 91–95, Nov. 2004, doi: 10.1109/MIS.2004.74.

[11] C. R. B. Azevedo, K. Raizer, and R. Souza, "A vision for human-machine mutual understanding, trust establishment, and collaboration," in *2017 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)*, Savannah, GA, USA, Mar. 2017, pp. 1–3, doi: 10.1109/COGSIMA.2017.7929606.

[12] A. Sciutti, M. Mara, V. Tagliasco, and G. Sandini, "Humanizing Human-Robot Interaction: On the Importance of Mutual Understanding," *IEEE Technology and Society Magazine*, vol. 37, no. 1, pp. 22–29, Mar. 2018, doi: 10.1109/MTS.2018.2795095.

[13] C. Vesper, "How to support action prediction: Evidence from human coordination tasks," in *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, Edinburgh, UK, Aug. 2014, pp. 655–659, doi: 10.1109/ROMAN.2014.6926327.

[14] L. M. Sacheli, E. Tidoni, E. F. Pavone, S. M. Aglioti, and M. Candidi, "Kinematics fingerprints of leader and follower role-taking during cooperative joint actions | SpringerLink," *Experimental Brain Research*, vol. 226, no. 4, pp. 473–486, 2013.

[15] I. Konvalinka, P. Vuust, A. Roepstorff, and C. D. Frith, "Follow you, Follow me: Continuous Mutual Prediction and Adaptation in Joint Tapping," *Quarterly Journal of Experimental Psychology*, vol. 63, no. 11, pp. 2220–2230, Nov. 2010, doi: 10.1080/17470218.2010.497843.

[16] A. Mörtl, M. Lawitzky, A. Kucukyilmaz, M. Sezgin, C. Basdogan, and S. Hirche, "The role of roles: Physical cooperation between humans and robots," *The International Journal of Robotics Research*, vol. 31, no. 13, pp. 1656–1674, Nov. 2012, doi: 10.1177/0278364912455366.

[17] P. Evrard and A. Kheddar, "Homotopy switching model for dyad haptic interaction in physical collaborative tasks," in *World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Salt Lake City, UT, USA, 2009, pp. 45–50, doi: 10.1109/WHC.2009.4810879.

[18] Y. Li, K. P. Tee, W. L. Chan, R. Yan, Y. Chua, and D. K. Limbu, "Continuous Role Adaptation for Human–Robot Shared Control," *IEEE Trans. Robot.*, vol. 31, no. 3, pp. 672–681, Jun. 2015, doi: 10.1109/TRO.2015.2419873.

[19] J. W. Crandall *et al.*, "Cooperating with machines," *Nat Commun*, vol. 9, no. 1, p. 233, Dec. 2018, doi: 10.1038/s41467-017-02597-8.

[20] F. Kaplan, "Talking aibo : First experimentation of verbal interactions with an autonomous four-legged robot," in *Learning to Behave: Interacting agents CELETWENTE Workshop on Language Technology*, 2000, pp. 57–63.

[21] Y. Mohammad and Toyoaki Nishida, "Human adaptation to a miniature robot: Precursors of mutual adaptation," in *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, Aug. 2008, pp. 124–129, doi: 10.1109/ROMAN.2008.4600654.

[22] S. Nikolaidis, D. Hsu, and S. S. Srinivasa, "Human-robot mutual adaptation in collaborative tasks: Models and experiments," *I. J. Robotics Res.*, vol. 36, pp. 618–634, 2017, doi: 10.1177/0278364917690593.

[23] F. Santoni de Sio and J. van den Hoven, "Meaningful Human Control over Autonomous Systems: A Philosophical Account," *Frontiers in Robotics and AI*, vol. 5, no. 15, pp. 1–14, 2018, doi: 10.3389/frobt.2018.00015.

[24] R. Crootof, "A Meaningful Floor for 'Meaningful Human Control,'" *Temple International and Comparative Law Journal*, vol. 30, no. 1, pp. 53–62, 2016.

[25] S. C. Calvert, D. D. Heikoop, G. Mecacci, and B. van Arem, "A human centric framework for the analysis of automated driving systems based on meaningful human control," *Theoretical Issues in Ergonomics Science*, pp. 1–29, Dec. 2019, doi: 10.1080/1463922X.2019.1697390.

[26] F. Ficuciello, G. Tamburrini, A. Arezzo, L. Villani, and B. Siciliano, "Autonomy in surgical robots and its meaningful human control," *Paladyn, Journal of Behavioral Robotics*, vol. 10, no. 1, pp. 30–43, Jan. 2019, doi: 10.1515/pjbr-2019-0002.

[27] G. Hoffman, "Evaluating Fluency in Human–Robot Collaboration," *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 3, pp. 209–218, Jun. 2019, doi: 10.1109/THMS.2019.2904558.

[28] L. P. J. J. Noldus, *The Observer XT*. 2019.

[29] E. M. van Zoelen, "Mutual Understanding for Human-Machine Collaboration: Evolving Leader-Follower Behavior," Master Thesis, Eindhoven University of Technology, Eindhoven, 2020.